

Military

EMBEDDED SYSTEMS

The COTS Technology Authority

EDITORIAL

> SDR: We need it *now*

INDUSTRY ANALYSIS

- > Product warranties
- > Military computing trends
- > "Beep, configure, fire"

OCTOBER 2005

VOLUME 1 NUMBER 2

www.mil-embedded.com

Mil Comms = SDR + FPGAs + DSP

SOFTWARE

SCA's evolution links
to other specs

HARDWARE


Altera logic
secures mil systems

APPLICATIONS

Low-power
flash storage

PRODUCT GUIDE

SBCs for signal
processing



Need Military Gigabit Ethernet?

GigEXTREME™ Lightens Your Load



ET-71000 GigEXTREME™ TCP/IP Offload Engine

	GigEXTREME™	Other TOE Cards	NICs with Partial Offload
Processor Loading*	5%	5%	50% to 70%
RTOS Support	✓	✓	—
No Application Modification or Recompile Needed	✓	—	✓
No Use of Upscreened Commercial Components	✓	—	—
Military Rugged	✓	Some	Some
IPv6 Capable	✓	—	—

* Processor loading measured for the TCP/IP receive network task on a 1 GHz Power PC host.



www.ddc-web.com

Italy +39 02 481 7900
Sweden/Denmark +46 8 544 90044
Benelux +31 78 621 5900
Finland +358 9 8194 220

Switzerland
Spain
Norway
Israel

DDC Europe

DDC (United Kingdom) LTD.
England
Tel: +44 (1635) 811140

DDC Electronique
France
Tel: +33 1 41 16 34 24
+41 41 748 32 32
+34 91 366 0159
+47 22917900
+972 9 796 6990

DDC Elektronik GmbH
Germany
Tel: +49 (0) 89 - 15 00 12 - 11

East Europe +44 1635 811140
South Africa +27 11 882 6836/7
Russia Federation +7 812 389 0104




ISO 9001 REGISTERED

ICS-572. Out of this world



Just because
you're short of space,
you shouldn't be
short on features.



The ICS-572 2-channel PMC module may be designed specifically for today's small-form factor, lightweight PCI systems, but that doesn't mean it compromises on functionality.

Combining the best in 14-bit ADC and DAC technologies with a high-density, high-speed, user-programmable FPGA, the ICS-572 delivers industry-leading cost, size and performance. It's ideally suited for multimission, high bandwidth radar and communications applications that need to both transmit and receive.

With support for Windows, Linux and VxWorks, and comprehensive application development support from ICS, the ICS-572 offers unparalleled flexibility and speed of deployment. As they say: small is beautiful.

ICS-572. Great leaps forward.



ICS

SENSOR PROCESSING

RSC# 3 @ www.mil-embedded.com/rsc

CHECK THE FULL SPEC NOW AT

www.ics-ltd.com/572

OR CALL TOLL FREE (US ONLY)

800-267-9794 or 613-749-9241

DEPARTMENTS

Industry Analysis

- 8 What has happened to product warranties?
By Jerry Gipper
- 10 Beep, configure, fire
By Don Dingee
- 12 Trends in military computing
By Joe Pavlat

In the System

- 16 FPGAs gearing up to dominate DSP applications
By Duncan Young, SBS Technologies

24 Editor's Choice Products

Crosshairs Editorial

- 62 We need SDR now
By Chris A. Ciuffo

56 New Products

By Sharon Schnakenburg

Product Selection Guide

- 50 Single board computers for DSP and reconfigurable computing

EVENTS

SDR Forum

Nov. 14-18 • Garden Grove, CA
www.sdrforum.org

Embedded Connect

Nov. 17 • Santa Clara, CA
www.esconline.com

COVER

A Humvee providing communication sits atop a lonely hill half covered in desert pattern camouflaging. (Photo courtesy of U.S. Marines by Lance Cpl. Jem J. Hamilton).

PRODUCT

Using internal PC/104-Plus modules, the Talon 8400 by GE Fanuc Embedded Systems is a self-contained PC designed for rugged applications from -20 °C to +50 °C.



Published by:



OpenSystems
Publishing™

©2005 Military Embedded Systems

FEATURES

HARDWARE:

20 Why convert to a SAASM-based GPS?

By Ron Holm, Symmetricom

25 Enabling military design security with high-performance FPGAs

By Joel Seely & Jie Feng, Altera

SOFTWARE: Software Defined Radio

28 Software communication architecture: Evolution and status update

By Jeff Smith, Ph.D., David Murotake, Ph.D., and Antonio Martin, SCA Technica

34 Cognitive radios: The future of SDR technology

By Bruce Fette, Ph.D., General Dynamics C4 Systems

38 Mapping waveforms to systems: What would a wideband networking waveform system require?

By Kevin Maier, Spectrum Signal Processing

APPLICATION: Performance enhancements

42 FPGA memory controllers improve DSP performance

By Richard M. Matthews, Micro Memory

46 Tapping into solid-state storage benefits for low-power and small form factor applications

By Gary Drossel, Silicon Systems

E-LETTER

WWW.MIL-EMBEDDED.COM/ELETTER

Net-centric computing: Architecting a distributed data-management infrastructure

By Bert Farabaugh, RTI

FLOSS helps you take control of your bytes

By Robert Dewar, AdaCore

Please visit us online for the complete list of articles

WEB RESOURCES

Subscribe to the magazine or E-letter:

www.opensystems-publishing.com/subscriptions



**"Great in a lab,
so how about
in a UAV?"**

Motorola makes products that are compute dense and famous for longevity like VMEbus and CompactPCI® boards and chassis using various PowerPC® and Intel® processors. They're produced commercially in significant volumes, so they're more competitively priced. How do you take full advantage of that in a field application?

**COTS by Motorola.
Mods by Arrow.**

That's where Arrow comes in. We have one of the broadest line cards in the industry and the expertise to integrate disparate systems, hardware, and software with extreme efficiency. We've gotten it right for thousands of clients for more than 29 years.

And Arrow teams with companies like Carlo Gavazzi Mupac, recognized across the industry for packaging and ruggedizing to meet military specs.

Reliability. Functionality. Flexibility. It's like getting full customization for a fraction of the cost. Yes. That'll fly.



MOTOROLA
Authorized Distributor

OEM Computing Solutions
888-427-2250
www.arrownacp.com

RSC# 5 @ www.mil-embedded.com/rsc

Registered trademarks are the property of their respective owners. They support systems.™ is a registered trademark of Arrow Electronics, Inc. Motorola and the Motorola logo are registered to the U.S. Patent & Trademark Office. CompactPCI is a registered trademark of the PCI Industrial Computer Manufacturers Group. PowerPC is a trademark of IBM Corporation and programs therein may be registered trademarks of International Business Machines Corporation or its subsidiaries in the U.S. and other countries.

Your support system™

ARROW

OpenSystems Publishing

Advertising/Business office:
30233 Jefferson Avenue
St. Clair Shores, MI 48082
Tel: 586-415-6500 ■ Fax: 586-415-4882

BUSINESS OFFICE:

Vice President Marketing & Sales
Patrick Hopper
phopper@opensystems-publishing.com

Business Manager
Karen Layman

ADVERTISING GROUP:

Vice President Marketing & Sales
Patrick Hopper
phopper@opensystems-publishing.com

Senior Account Manager
Dennis Doyle
ddoyle@opensystems-publishing.com

Account Manager
Tom Varcie
tvarcie@opensystems-publishing.com

Print and Online Marketing Specialist
Christine Long
clong@opensystems-publishing.com

Advertising/Marketing Coordinator
Andrea Stabile
astabile@opensystems-publishing.com

European Bureau Chief
Stefan Baginski
sbaginski@opensystems-publishing.com

Account Manager
Doug Cordier
dcordier@opensystems-publishing.com

COMMUNICATIONS GROUP:

Patrick Hopper
phopper@opensystems-publishing.com

Christine Long
clong@opensystems-publishing.com

EMBEDDED GROUP:

Dennis Doyle
ddoyle@opensystems-publishing.com

Doug Cordier
dcordier@opensystems-publishing.com

MILITARY GROUP:

Tom Varcie
tvarcie@opensystems-publishing.com

Andrea Stabile
astabile@opensystems-publishing.com

INTERNATIONAL SALES:

Stefan Baginski
sbaginski@opensystems-publishing.com

For reprints and PDFs call the sales office: 586-415-6500

Military

EMBEDDED SYSTEMS

AN OPENSYS TEMS PUBLICATION

Military and Aerospace Group

- DSP&FPGA Product Resource Guide
- DSP-FPGA.com
- DSP-FPGA.com E-letter
- Military Embedded Systems
- Military Embedded Systems E-letter
- PC/104 Embedded Solutions
- PC/104 Embedded Solutions E-letter
- PC/104 & Small Form Factor Catalog
- VMEbus Systems
- VMEbus Systems E-letter

Group Editorial Director Chris A. Ciufu
cciufo@opensystems-publishing.com

Managing Editor Bonnie Crutcher
bicrutcher@opensystems-publishing.com

Senior Editor (columns) Terri Thorson
tthorson@opensystems-publishing.com

Assistant Editor Sharon Schnakenburg
sschnakenburg@opensystems-publishing.com

European Representative Hermann Strass
hstrass@opensystems-publishing.com

Art Director Steph Sweet

Senior Web Developer Konrad Witte

Graphic Specialist David Diomede

Circulation/Office Manager Phyllis Thompson
subscriptions@opensystems-publishing.com



Editorial/Production office:
16872 E. Ave of the Fountains, Ste 203, Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466
Website: www.opensystems-publishing.com

Publishers John Black, Michael Hopper, Wayne Kristoff

Vice President Editorial Rosemary Kristoff

Communications Group

Editorial Director Joe Pavlat
Assistant Managing Editor Anne Fisher
Senior Editor (columns) Terri Thorson
Technology Editor Curt Schwaderer
European Representative Hermann Strass

Embedded and Test & Analysis Group

Editorial Director Jerry Gipper
Editorial Director Don Dingee
Associate Editor Jennifer Hesse
Senior Editor (articles) Terri Thorson
Technical Editor Chad Lumsden
Special Projects Editor Bob Stasonis
European Representative Hermann Strass

Canada return address: WDS, Station A, PO Box 54, Windsor, ON N9A 615

Military Embedded Systems is published six times a year by OpenSystems Publishing LLC., 30233 Jefferson Ave., St. Clair Shores, MI 48082. Subscriptions are free, upon request in writing, to persons dealing with or considering *Military Embedded Systems*.

For others inside the US and Canada, subscriptions are \$34/year. For 1st class delivery outside the US and Canada, subscriptions are \$60/year (advance payment in US funds required).

POSTMASTER: Send address changes to *Military Embedded Systems*
16872 E. Ave of the Fountains, Ste 203, Fountain Hills, AZ 85268

Rad-Tolerant. Space-Approved. COTS Compliant.

And Designed, Built and Tested by Aitech.

Designing and building board level products and sub-systems for space applications is tough enough. Doing it in a true COTS environment is even tougher. It takes a very special company to do it right – and that company is Aitech. We not only designed and built the world's first harsh environment, open architecture VMEbus boards more than two decades ago, but we're fully qualified for use in today's most hostile environment – space.

The only COTS company... Aitech is the only COTS company in the world that offers embedded products for space applications with this combination of features: •Designed and qualified specifically by us for space •Radiation-tolerant throughout •On-board triple redundant memory •Rad-hard SOI (silicon on insulator) ASICs •Single event effects and total dose radiation survivability


Total space applicability... Aitech embedded products and sub-systems for space are ideal for Low, Medium and High Earth Orbit applications, lunar and Mars robotic vehicles and much more. Our products are used in the Space Shuttle, MIR Space Station, International Space Station and other high profile satellite programs where highest performance and reliability are required.

Real space-qualified COTS... not custom off-the-shelf...but commercial off-the-shelf. COTS the way it's supposed to be!

We don't compete with you... some embedded companies try to be systems integrators. We don't. We deliver board-level product and integrated sub-systems for space (and military/aerospace) applications. We leave the systems integration to the companies that do it best...like yours!

We have what you need... from a full range of high performance, low cost, rad-tolerant, space-qualified CompactPCI SBCs, peripheral I/O boards and PMCs, memory boards and enclosures...to complete radiation and qual-testing, component obsolescence risk mitigation, lifecycle support and program management capabilities – with all the economy-of-scale advantages of off-the-shelf products.

Make us prove it... we can and we will. Call or visit us on the web. Embedded space is *our* space.



Aitech Defense Systems, Inc.

9301 Oakdale Avenue
Chatsworth, CA 91311

email: sales@rugged.com

Toll Free: 888-Aitech8 - (888) 248-3248

Fax: (818) 718-9787

www.rugged.com

What has happened to product warranties?

By Jerry Gipper



Most embedded electronic hardware components come with a standard product warranty. If a board fails, you contact the supplier, arrange a repair or exchange, ship it to the repair center, and await the repaired or replaced product. This is a rather routine process amongst all suppliers worldwide. In the case of most VME, CompactPCI, and embedded motherboards, today's warranties range from one to three years in duration.

There used to be a time when warranties from many suppliers were five years or in some cases, "lifetime" (Figure 1). What ever happened to those warranties? Why are they not as common? In the early 1990s there was a lot of noise in the industry about long or lifetime warranties. Press releases were distributed touting the warranty terms, and warranties were highlighted in advertisements. Suppliers were very proud of their product reliability and were willing to stand behind their products.

A quick check today of the same product types shows typical warranty periods of only two years, with a few as long as three years. Certainly these are shorter than "lifetime." Has something happened to the reliability? Has the cost of quality risen too high to absorb in competitive pricing strategies?

This change in warranty periods started quietly happening as companies were recovering from the dot-com bust in 2001. No embedded technology suppliers publicly "announced" their shorter warranty periods. In fact, the reduced warranty periods did not hit home until a problem occurred and a product needed repair. Did suppliers reduce prices to reflect the shortened warranty periods? I would venture to guess prices were not reduced since suppliers were already under extreme margin pressures during those times.

Some possible explanations exist as to why warranties greater than two years are no longer typical. Any one or more of the

following are key factors in determining warranty terms and pricing:

- **Increased cost of quality:** By definition, the Cost Of Poor Quality (COPQ) consists of costs generated as a result of producing defective material. Additional cost includes all the labor cost, rework cost, disposition cost, utilities cost, manufacturing cost, and material cost that has been added to the unit up to the point of rejection. It also includes the cost of lost opportunity due to the loss of resources used in rectifying the defect. The cost of lost opportunity means lower revenue and profit, potential loss of market share, and a lower service level to customers. All areas of cost in COPQ have increased in past years.
- **Speed of technology:** In general, the technologies used in embedded computing systems are still relatively new. New technologies evolve rapidly, and rapidly evolving technologies have a shorter product life cycle. Getting the quality to an acceptable level can be a challenge at best with many of the newer technologies. Do designers and manufacturers always fully understand the quality issues of one technology before moving on to the next?
- **Pursuit of quality:** With products' shorter life cycles, perhaps the race to have the best and most optimized set of features at a competitive price has pushed quality into the backseat. Products today have manufacturing lives less than three years in duration. In most cases, the quality is not fully optimized in that short of a life cycle. Suppliers move on to the next generation before they really get the kinks out of the existing products. Sometimes designers may take shortcuts or use unproven design elements to meet time-to-market pressures. With the low quantities and high mix of the embedded space,




Figure 1

it is hard to spend too much time on quality for any single product line. Is enough simulation and accelerated life testing being done?

- **Lost the recipe:** Perhaps manufacturers simply “lost the recipe” when changing manufacturing locations as mergers occurred or manufacturing was transferred offshore to large, high-volume contract manufacturers to save costs. The embedded industry tends to

build products in small lots with a high mix of different types and product versions. The manufacturing efficiency is never very high in these situations. Moving to a high-volume, low mix manufacturing facility only frustrates the manufacturing process in these cases. As a result, quality suffers immensely. Lower quality levels mean higher warranty costs and loss of ability to absorb these costs in competitive prices. Most manufacturing operations are great at


building high-quality products, and vendors fully understand the nuances of quality manufacturing; however, it takes high repetitions of the same steps and processes to really reach the best quality levels.

- **Service as a profit center:** Maybe the answer is as simple as using warranties as a way to increase the revenue stream. Most embedded board suppliers have extensive repair and service centers. Removing the cost of the warranty from the product cost and making it a separate price line item is a way for the service centers to improve their revenue streams and justify their existence. It does make it easier to develop custom and innovative service programs that better fit the customer's needs.

What is going to happen to warranty terms in the future? It is possible that the Restriction of the use of certain Hazardous Substances (RoHS) or the Waste from Electric and Electronic Equipment (WEEE) initiatives being driven out of Europe and Asia will force further tightening of warranty terms. Suppliers may use this as the compelling event to tighten the warranty programs even more. Two-year warranties could easily become one-year warranties in the next one to two years.

However, all is not lost as many suppliers in the embedded industry offer various extended warranty programs that can be customized to meet the specific needs of consumers. These programs have options for warranty periods and coverage that can help manage costs for the lifetime of a program. Extended warranties are insurance against a future failure. Be sure to check out all your options when planning lifetime costs for your products.

For more information, you may contact Jerry directly at jgipper@opensystems-publishing.com.



Proven COTS Solutions from
the FPGA Computing Experts



VME, cPCI, PCI and PCI-104 COTS Hardware

Systems Software and Development Tools

To find out more and register for the
latest Virtex-4 product information visit -
www.nallatech.com/capability

North America Toll free: 1-877-44-NALLA
EMEA & ROW Phone: +44 (0) 1236 789500

 **NALLATECH**
The High Performance FPGA Solutions Company

Beep, configure, fire

By Don Dingee



Military planners are now talking in terms of *plug-and-fight* capability. This is a spin on the old personal computing term *plug-and-play* – catchy, but outdated. As a marketing person at heart, I can't help but suggest an updated term that better describes what we see taking place and what elements are going into this capability.

Most of us have seen some version of the movie scene where the condemned prisoners are paraded out in blindfolds, lined before a wall, and a firing squad is issued the familiar cliché "Ready, aim, fire!"

Today, with electronics driving most combat systems, a more appropriate term for the plug-and-fight process would be "Beep, configure, fire!"

Beep – Join the network

For today's warfighter, there is a wealth of information and intelligence available in real time as close as the nearest network. Sensor and command data fused into the network from a variety of modes and locations enable a tactical system to quickly access reliable information on the mission plan, threat situation, and the configuration of neighboring defense systems.

Recently, I visited a friend in the San Francisco Bay area, and he was showing me his new Bluetooth-enabled phone and headset. He commented that he could not help but notice that when he wears the headset while driving, especially on the freeways, he hears frequent beeps as the headset attempts to connect with devices in neighboring vehicles. It's almost *too* easy to connect to a network now.

Short-range wireless networks provide a continuous network tone and have made it easy to connect quickly to a wide variety of devices. Technologies such as Bluetooth, Wi-Fi, Ultrawideband, Wireless USB, and the longer-range WiMax may soon be enabled with a security layer, enabling their use in defense applications. For now, defense networks using Software Defined Radio (SDR) technologies fit the bill.

Planners talk of systems fitting into a *netted, distributed* force. Whether wired or wireless, today's defense systems-of-systems are built around network-centric designs and rely on being able to join the network quickly.

Configure – Compute for effect

Reconfigurable computing has advanced an incredibly long way in recent years. Advances in System-on-Chip (SoC) technology combining processing with FPGAs allow designers to optimize

and adapt computing hardware to better fit immediate requirements. Linux-based open system architectures allow *clusters* (just another term for systems-of-systems) to add distributed nodes easily, enabling distributed processing and sensor fusion capability quickly and reliably.

It's the intersection of processor, network fabric, and operating system in a reconfigurable form that becomes interesting. One example of this is the CoSine from Micro Memory. Built on the Xilinx Virtex-4 FPGA family with VxWorks running on two PowerPC 405 cores, it includes connectivity for serial RapidIO or PCI Express, a corner-turning Direct Memory Access (DMA) engine to aid in data movement, and a user-programmable area for signal processing algorithms. (Refer to the Micro Memory article on page 42 in this issue.)

"Today, with electronics driving most combat systems, a more appropriate term for the plug-and-fight process would be 'Beep, configure, fire!'"

Additionally, advances are being made in Linux on FPGA platforms. The University of Queensland in Australia recently secured a contract for NASA's Reconfigurable Scalable Computing (RSC) project, building from a starting point in which a version of uClinux is adapted to the Xilinx MicroBlaze soft-core processor. The ultimate goal will be to remotely (as in ground-to-space) partially reprogram FPGAs on the fly.

We may see that general-purpose processors ultimately give way to reconfigurable processors in defense applications, both for functional and life-cycle reasons. Teamed with the University of Southern California and using IBM Cu-08 90 nm process technology, Raytheon's Morphable Networked Micro-Architecture (MONARCH) project targets April 2008 for a device that can alternate instantaneously between front-end (streaming) and back-end (threaded) processing.

According to Jack Kelble, president of Raytheon Space and Airborne Systems, "In the past, a bank of processor boards accepted information and another bank processed it." He added, "Now a tiny but highly sophisticated device a fraction of the size will perform both functions with unprecedented speed, power, and capacity to store and process a vast amount of data." MONARCH is expected to perform in a single chip or SoC role, possibly significantly reducing the number and types of processors required for computing systems. In the latter application, it will be able to process incoming and outgoing data while analyzing it.

Fire – With open weapons

Using technologies such as these, networked, reconfigurable weapons systems are beginning to emerge as the norm. Initiatives including the Modular Open Systems Approach (MOSA) spon-

sored by The Open Systems Joint Task Force within the US Department of Defense continue to set higher expectations.

Openness in new acquisitions is being demanded as seen in these comments from Lieutenant General Ronald E. Keys in February 2005: "We've got to get to this thing called the 'compatible open architecture.' I've got to be able to truly plug-and-play, and it's got to plug-and-play better than Microsoft. It's actually got to plug, boot up, recognize, and work ... So don't bring me stuff that's not compatible because I'm not going to be happy."

Large combat systems already architecting around plug-and-fight concepts include the Medium Extended Air Defense System (MEADS), Future Combat Systems (FCS), and the Littoral Combat Ship (LCS).

MEADS is the US Army's next-generation replacement for Nike Hercules, Hawk, and Patriot air-defense missile systems, designed from the ground up to move with ground forces and interoperate with other allied forces. It relies heavily on networking and distributed intelligence to achieve its mission. A MEADS system has the capability to command a fleet of distributed missile launchers while simultaneously detecting and tracking hostile forces and targets. There is a key tactical advantage to this distributed design: The missile launchers can be located well away from the ground radar and the battle management units, reducing the risk of detection of the launchers. This tactical advantage also opens the possibility to transfer command and control of the launchers and missiles to a neighboring battle management unit, while some management systems are offline for whatever reason.

FCS isn't a single system but rather a blended system-of-systems intended to transform the US Army's fighting capability. Underlying FCS is a software architecture called Fire Control – Node Engagement Technology (FC-NET). FC-NET provides an adaptable, flexible architecture that modularizes the interaction between the technical weapon system (the intelligence that controls and guides the weapon) and the tactical information systems. This enables weapon systems to readily join the command fabric to get the information they require.

Another plug-and-fight system is LCS. It's being designed to work in three primary mission areas for the US Navy, including mine countermeasures, anti-submarine warfare, and anti-surface warfare, presumably with anti-air, self-defense capability in each role. This is being accomplished through the design of mission packages that fit into the sea frame and adapt the capability to the desired mission area. Open-systems architecture and modular, networked subsystems are again the key to success, and the notion of being able to reconfigure the system for the role at hand is prominent in the architecture.

Open doesn't mean big

Creating new systems-of-systems isn't necessarily about using big computers. From the looks of things, it could be just the opposite, using networks of relatively small processors tied together wirelessly with very intelligent software and combining these systems into larger systems.

Researchers at the University of Essex are working on a concept called the *gridswarm*, where small Unmanned Aerial Vehicles (UAVs) capable of speeds up to 120 mph fly in formations similar to the flocking behavior of small birds. In the prototype, these aircraft are connected by a Bluetooth mesh driven by Linux compute modules from Gumstix. These tiny modules run Linux 2.6 on 400 MHz Intel Xscale processors with 64 MB DRAM and 4 MB Flash, along with USB, serial, and optional Bluetooth interfaces. It's a great example of small systems fitting into larger systems fitting into still larger systems with aggregated intelligence.

Rapid developments in wireless networking, reconfigurable computing, and network-centric weapons systems are going to spawn new innovations quickly. The results should also reduce the long-term costs of weapons procurement, enabling easier upgrades and reducing the impact of obsolescence by allowing subsystem level replacements.

I'll be sure to tell my friend the next time I see him that when he hears a whole bunch of beeps in rapid succession on his Bluetooth headset, he should duck. It could be a UAV gridswarm reconfiguring just overhead, and hopefully they are unarmed and peace loving.

If you happen to see a gridswarm, or other interesting developments that beep and configure, drop me a line.

For more information, contact Don at ddingee@opensystems-publishing.com.

**All Avionic Protocols
in the Palm of your Hand**

ISA • PCI • PC104 • cPCI • PCM • VME
MIL-STD-810 • Custom Boards

BMC Communications Corp.
67 Bond Street
Westbury NY 11590
Tel: (516) 997-2118
Fax: (516) 997-2129
UADI@BMCcorp.com

BMC
www.BMCcorp.com

MIL-STD-1553 ARINC 453/708 ARINC 429, 575, 561, 571, 572, ... RS232/422

RSC# 11 @www.mil-embedded.com/rsc

Trends in military computing

By Joe Pavlat



A changing landscape

These are good times for designers of military computer systems, as the range of choices and the breadth of applications and requirements is greater than ever before. Traditionally, military electronics have been extremely expensive, usually purpose-designed and uniquely built for each application. Systems often do not communicate with each other, making future net centric warfare difficult. Reuse of hardware and software has been the exception rather than the norm, and design cycles historically have been long and expensive. Industry insiders often talk about the large “flywheel” in the military computer business, meaning that development cycles are long, and revenues are often years away. The one open standard embraced for military applications, the VMEbus, arguably has been widely accepted not for its blazing performance but rather because the standard was first published more than 20 years ago – a veritable lifetime for the rest of the computer industry.

The flywheel is still quite large, but it is rotating a bit more quickly these days. Open industry standards are becoming more popular. Former Secretary

of Defense William Perry’s famous COTS directive was a factor, but the same development cost and time-to-deployment pressures that affect the commercial computer world affect military suppliers. Two PCI Industrial Computer Manufacturers Group (PICMG) standards, PICMG 1.0 (PCI-ISA Passive Backplane, 1994) and PICMG 2.0 (CompactPCI, 1995), were released 11 and 10 years ago respectively, and are used for a variety of military applications worldwide.

A wide range of open standards

PICMG’s first published specification, the PCI-ISA Passive Backplane Specification released in 1994, is being used in a wide range of applications, including the on-board fire control computer for the M109 Paladin self-propelled howitzer used extensively in Operation Iraqi Freedom. Companies such as BES Systems Ltd. in Israel offer a complete range of ruggedized airborne, vehicle, and naval computers compliant to the PICMG 1.1 specification, additionally providing compliance to military standards including MIL-STD-810E, which dictates tough requirements for shock, vibration, humidity, fungus, salt and dust, and fog.

Released in 1995, the CompactPCI standard was developed for ruggedized industrial applications. It offered then state-of-the-art performance, based on ubiquitous PCI silicon available from virtually every microprocessor and peripheral chip manufacturer. It was based on the same IEEE 1101.1 mechanical standard used by VME, and it became very popular for communications applications worldwide. Defined for both 3U and 6U form factors, the 6U size became popular for the vast majority of communications applications, which needed every square inch of real estate for components. The 3U form factor has historically been used largely for instrumentation and some industrial automation applications, but was not as widely embraced as the larger 6U form factor.

This has been changing over the last few years in a dramatic fashion. Specially ruggedized 3U CompactPCI products are being used for a wide variety of airborne, vehicle, and even space-based systems. One example is the AVC-CPCI 3009 system offered by SBS Technologies, developed for Unmanned Aerial Vehicle (UAV) applications (see Figure 1, photo courtesy of SBS Technologies, Inc.). Its integrated frame grabber and MPEG-4 image compressor connect directly to the airframe’s onboard camera, forwarding data in real time to war planners on the ground.

Systems are also going into space. Aitech’s S950 3U CompactPCI SBC is conduction-cooled and offers a PowerPC 750FX CPU (see Figure 2, photo courtesy of Aitech Defense Systems, Inc.). It is rated to operate in Low Earth Orbit, Geosynchronous Orbit, and Mars Terrestrial environments.

The 6U CompactPCI systems are also being used for military applications. Performance Technologies, Inc. builds a sophisticated Mission LAN System using the PICMG 2.16 CompactPCI Packet Switched Backplane standard. Intended to be part of a National



Figure 1



Figure 2

Command Center aboard a heavily modified Boeing 707 aircraft known as the *TACAMO*, the system maintains communication and control in the event that other command centers are damaged or destroyed. It provides networking and routing within the aircraft, handling packetized radio, satellite, radar, and laser transmissions, and ties together different systems on the plane (see Figure 3, photo courtesy of Performance Technologies, Inc.).

Performance Technologies has also developed a unique hybrid CompactPCI/

VME system for use in the Global Hawk UAV (see Figure 4, photo courtesy of Performance Technologies, Inc.). This computer provides near real-time high-resolution images and intelligence to field commanders in theater or across the world. Multicast image streams can be ordered by a commander in a control room or a soldier on the ground in the next valley, providing vital current information and situational awareness. The CompactPCI boards are conduction cooled and compliant with the ANSI/VITA 30.1 specification (2 mm connector practice for conduction-cooled Eurocard systems).

Better architectures hit the ground

One hears a great deal about transforming the military from the current platform-centric approach to network-centric operations. The underlying computer technologies, including chips and software, are also undergoing fundamental change, and it is good news for designers and users of military computers.

Most backplane interconnect technologies, including VME and CompactPCI, are based on chip-level interconnects that were intended for planar motherboards. Hot swap was not an integral part of these interconnects, and their parallel nature has meant that any board plugged into the backplane could cause the entire system to fail if it failed. Full 2N redundancy was often the only solution. Focus was placed on *reliability* instead of the much more useful concept of *availability* because parallel bus architectures just do not adapt well to high availability designs, which require system management and failure domains of a single board. Also, as core chip voltages go ever lower, the notion of distributing chip supply voltages often means that parallel backplanes are required to produce hundreds, or even thousands, of amps of current for large systems.

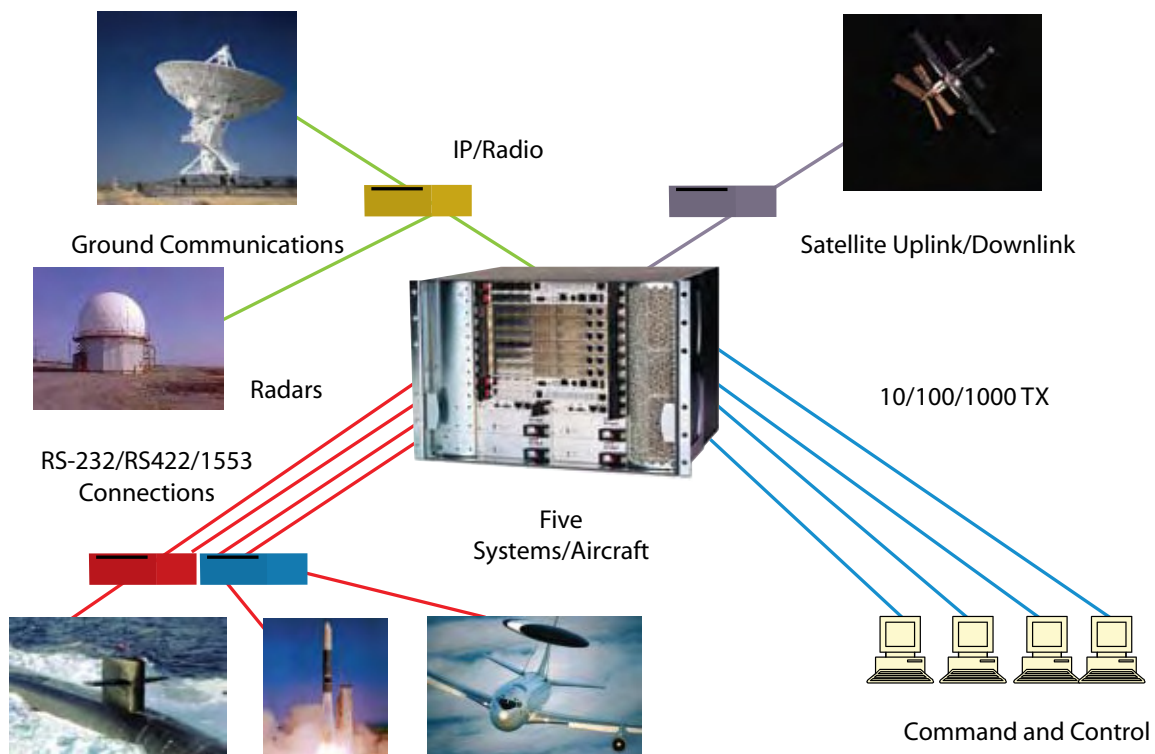


Figure 3




Figure 4

This is changing for the better. Primarily, in order to increase data transfer speed, the microprocessor industry is moving rapidly towards switched serial interconnect technology, popularly known as *switch fabrics*. This technology reduces the speed-robbing capacitance typical of a parallel bus architecture with instantaneous point-to-point interconnects. Not only do these interconnects increase data transmission speed one or two orders of magnitude, which is important for military applications such as imaging, they can, if properly designed, reduce the failure domain to a single board or Field Replaceable Unit (FRU). They also dramatically improve the scalability of military systems, as the same packetized data format used over the backplane can be used between boxes or systems in a large network. Switched fabric Ethernet-based backplane standards, first introduced to

the world in 2001 in the PICMG 2.16 Specification, are beginning to be used for military applications.

Additionally, power distribution concepts are changing with an emphasis toward shipping higher voltages across backplanes in order to reduce the ever-increasing currents required by Moore's Law. (We'll save cooling problems for another time.) Localized power conversion is the norm in standards like AdvancedTCA and PICMG's recently ratified advanced mezzanine card standards.

An entire book could be written about software development, but as much as advances in this arena seem to trail hardware progress, a few things can be said. Spurred by the wide adoption of the PCI bus 10 years ago, military board suppliers are increasingly being freed from the need

to write and maintain cumbersome Board Support Packages (BSPs), which often need updating every time a chip on the board changes revision. The Application Programming Interface (API) approach moves that responsibility to the chip supplier and operating system, making systems easier to upgrade and maintain. Expensive RTOSs are beginning to give way to less expensive and increasingly powerful OSs such as Real-Time Linux and Carrier-Grade Linux.

Notions are changing

In the commercial communications sector, the distinction between datacom and telecom is all but gone as the world's infrastructure moves towards packet-based communications. Military infrastructure, at least in the US, is joining the movement. Major initiatives, such as the Department of Defense's Warfighter Information Network Tactical (WIN-T) program, are based on commercial communication technologies, including secure wireless networks, Voice over Internet Protocol (VoIP), PCS cellular services, and ATM data transport. PDAs, laptops, and tablet computers are widely used in American command centers worldwide, and e-mail is as ubiquitous and important as it is in the civilian sector. The old notions about ruggedized military computers being completely customized boxes milled out of large bars of aluminum are changing. And they are changing for the better as the flywheel spins faster.

For more information, contact Joe at jpavlat@opensystems-publishing.com.

Crystal P4 875



- Dual 10/100/1000 network Ethernet
- Up to 800 MHz FSB supporting 128k through 2MB L2 Cache
- PICMG 1.2 with 64 bit/66 MHz support
- Two SATA, One Ultra ATA/100, Floppy support

For more information, contact a Program Manager at:

crystal
GROUP INC.

Hiawatha, Iowa USA 52233-1204
1.800.378.1636 1.319.378.1636
www.crystalpc.com

Danville Signal

Embedded SHARC dspblok



DSPBLOK DSP MODULES

Danville's dspblok™ DSP Function Modules are small (60 x 60 mm), highly integrated SHARC DSP cores that are easy to incorporate into embedded applications. Each dspblok includes a DSP, flash, eeprom and a switching power supply brought out to 2mm headers. A dspblok may also include a user programmable FPGA or PLD and SDRAM.

www.danvillesignal.com

COTS modular DSP systems:

scalable
flexible
real-time &
high performance

- PCI
- cPCI
- USB
- standalone

FPGA &/or DSP &/or I/O

tools, IP tutorials included

Supporting Your Future
HUNT ENGINEERING

www.hunt-dsp.com

Acromag I/O. Customize It.



Acromag introduces affordable FPGA I/O. For ALL your projects.

As an engineer, your projects are unique, ever-changing, and budget-bound. That's why our new PMC modules give you an affordable solution to create custom I/O boards.

But if you thought FPGA computing was only for top-end applications, think again. Our PMCs are ideal for protocol conversion, simulation, in-circuit testing, and much more. So why settle for generic I/O when you can design exactly what you need while staying in budget and reducing your time to market?

- Virtex®-II FPGA with 500K, 2M or 3M system gates
- 1Mb on-chip RAM, 9Mb on-board SRAM
- Fast PCI with 32-bit, 66MHz dual DMA



Industry Pack FPGA I/O also available

Cost-effective custom I/O

Choose from a variety of I/O configurations:

- Digital I/O: TTL, RS422, or LVDS I/O
- Analog I/O: four 20 or 65MHz A/D and two D/A

Faster time to market

Why waste precious time building a board from scratch? Our new FPGA modules let you process your I/O signals any way you want. Quickly.

Flexibility to meet unexpected challenges

Acromag FPGA I/O will help you bring your projects in on time and under budget. And with FPGAs, you'll be ready to adapt to all the inevitable changes.

Thinking about FPGA I/O?

Think flexible. Think affordable.

Think Acromag.



Call or visit our website today – for VME, CompactPCI, PCI, PMC, and Industry Pack solutions.

Acromag 
THE LEADER IN INDUSTRIAL I/O

www.acromag.com

800-881-0268 or 248-624-1541

ANALOG I/O

DIGITAL I/O

SERIAL I/O

FPGA

COUNTER/TIMER

QUADRATURE

Manufactured in Wixom, Michigan, USA

RSC# 15 @ www.mil-embedded.com/rsc

All trademarks are the property of their respective companies

FPGAs gearing up to dominate DSP applications

By Duncan Young



AltiVec PowerPCs have enjoyed dominance in high-end military signal processing applications such as Software-Defined Radio (SDR) and video/image compression. But FPGAs, with their inherent parallelism and reprogrammability, are fast becoming prevalent and easily fit on smaller boards such as PCI Mezzanine Cards (PMCs).

The PowerPC's AltiVec, 128-bit vector processor is the undisputed first choice for DSP where floating-point performance is the critical decision point. However, the new breed of FPGA is geared up to wipe out the AltiVec advantage for the great variety of the military's fixed-point DSP applications. Interestingly, although the AltiVec has been available for some time and appears in Freescale Semiconductor's recently announced dual-core MPC8641D embedded processor, there does not appear to be any interest within the industry in offering competitive (and hence better or faster) floating point solutions for military DSP applications.

FPGAs with their innate parallelism appear set to oust the traditional multiple-PowerPC solution from a broad range of DSP applications. FPGAs offer much reduced real estate, power consumption, and cost. Time-to-deployment pressure from end users is driving some COTS vendors to offer prepackaged FPGA-based hardware/software products for specific application segments.

Let's talk numbers

The advent of the latest generation of FPGAs such as the Xilinx Virtex-4 or the Altera Stratix II — having as many as 100,000 logic cells, 400 DSP Multiply/Accumulate (MAC) elements and internal RAM, clocking at over 450 MHz — offers the potential to migrate and integrate DSP operations previously only performed by dedicated processors. The ability to process multiple data streams in parallel sets the FPGA solution apart from dedicated DSP or PowerPC alternatives.

Where ultimate performance is required on a single channel, the FPGA will always lose out. The AltiVec can perform a 64 x 64 MAC at single instruction rates (typically 1.5 GHz), whereas the FPGA will achieve 450 MHz for an 18 x 18 MAC. The FPGA's aggregate throughput from 400 MACs on multiple channels can far exceed that of a single PowerPC processor device in most applications. The DSP system designer can now choose from a number of design directions:

Front-end processing of multiple array sensors

Unlike a PowerPC or dedicated DSP processor, an FPGA offers many MAC elements to perform operations in parallel on incoming data streams. Previously this kind of parallelism required many PowerPCs to perform repetitive tasks such as filtering and decimation from each array of the sensor, with the attendant data distribution challenges. This data distribution is simple if all the PowerPCs are located on, for example, just a single VMEbus card, but much more complex if they are spread among many cards or even racks of cards where switched fabrics would be required. Using FPGAs can bring about big reductions at the critical front end, typically offering savings equivalent to one FPGA per 10x PowerPCs. Power, weight, space, complexity, and cost are conserved.

Complete system-on-chip fixed-point DSP solutions

The FPGA provides a versatile set of external interface options such as LVDS, PCML, LVPECL, and HyperTransport with which to build multiple types of sensor interfaces. The DSP and logic elements can be used for filtering and FFT operations, plus onboard hard and soft processor cores and real-time operating systems are available for logical and system-level operations. Extensive tool kits and libraries are available for the development of FPGA architecture, routing, and software to suit the final application requirement.

With the notable exceptions of multiple-array radar and sonar applications, most military applications for DSP are relatively straightforward and can be closely coupled to the sensor itself. Typical of this class of application are:

- Video and image processing
- Communications (SDR)
- Weapons system sensors such as torpedoes, cruise missiles, ground-to-air missiles, air-to-air missiles, or surveillance UAVs.

Unlike the multiple-array radar and sonar applications that often use many hundreds of PowerPC-class processors plus a switched fabric for interconnection, these simpler applications listed above could be implemented by dedicated DSP devices such as the TigerSHARC, ASICs, FPGAs, or by a much smaller number of PowerPC processors.

SDR example

A potential application for a single FPGA device could be to form the basis of a typical rugged, man-portable SDR with multiple secure voice and data channels. Soft FPGA cores are available for filtering, and IFFT/FFT modulation and demodulation using the FPGA's DSP elements, plus encryption and decryption — thus providing a number of communication channels in just one FPGA device. To offer complete, system-level functionality, this configuration needs to be augmented with external interfaces to the RF components and the user, plus a general-purpose processor for operation of the user interface and regular diagnostics/prognostics of the completed system. Required user interfaces are:

- Serial RS-232/422 channels
- Ethernet
- USB
- Discrete I/O

An ideal packaging solution for an FPGA front-end is the PMC module. Most generic SBCs from established military



Figure 1

COTS vendors support the PMC concept of I/O, as indeed do many PowerPC-based DSP cards. The PMC module is independent of the SBC's processor type or bus architecture (for example VMEbus or CompactPCI), making it a versatile platform for many different applications. In this example, if the FPGA were implemented on a PMC-format module, it could then be mounted on a COTS 3U CompactPCI host SBC with either a Pentium or PowerPC processor. This combination of SBC and PMC would form the digital portion of the radio and would only occupy a single 0.8" slot width, 6.3" deep and 4" high.

Time-to-deployment of new FPGA designs

Time-to-deployment of new technology has become a critical factor in the military procurement process. Whereas projects would once take many years to reach combat status, with what was by then obsolete technology, timescales have shrunk. Introducing new capabilities such as security and surveillance now demands not just the latest technology but all the tools and application support required for their immediate use.

Though SOC FPGAs promise the benefits of technology leadership, there is considerable development effort required to reach deployment when undertaking new application designs. This requirement is so even if it is based on a typical off-the-shelf solution such as a PMC module

using the FPGA manufacturer's toolkits. Some COTS vendors have taken the idea further and developed prepackaged applications for their FPGA. PMC modules that include standard hardware interfaces and application-ready code can be used

in areas such as SDR and video compression-based surveillance. Bundled COTS products reduce time-to-deployment to satisfy new expectations.

The TS-MPEG-4 bundled product from SBS Technologies shown in Figure 1 illustrates this product packaging strategy. It is based on a standard air- or conduction-cooled PMC module with an Altera Stratix EP1S30 FPGA, 128 MB of SDRAM, and a PCI interface to the host SBC. This PMC can be used in its basic form in many different applications with the customer developing code to suit. The external interfaces can be tailored by means of a unique micro-mezzanine mounted on the PMC module itself.

In the module's bundled form, the deliverable product includes all the code required for operation out of the box, plus physical interfaces to two RS-170 video sources and one RS-170 video output. Instead of using DSPs or general-purpose processors such as AltiVec-equipped PowerPCs, the supplied FPGA code compresses the two incoming video channels using MPEG-4, offering 15 to 20 percent

COTS


DATA BUS INTERFACE TRANSFORMERS

MIL-PRF-21038

SWITCH MODE POWER MAGNETICS

PREMIER MAGNETICS is your source for MIL-STD-1553 databus interface transformers for avionics, fly-by-wire and guidance systems. And we're your source for switch mode power supply components—high-frequency transformers, inductors and filter components—complete with fully-tested reference designs.

- Smallest Industry Package Sizes
- Competitive Prices
- Prompt Delivery
- Free Samples
- Reference Circuits



Premier Magnetics
 20381 Barents Sea Circle
 Lake Forest, CA 92630
 949-452-0511 fax: 949-452-0512
www.premiermag.com/mes

In the System

better compression ratios than the widely used MPEG-2 standard. The card then streams compressed data over the PCI bus to the host SBC. On the host, the MPEG-4 stream is encapsulated and passed to a network connection using UDP. Remote display clients supplied with the bundle can then decompress and display the decoded video streams. It should be noted that this is an FPGA-based system, reaping all the benefits of FPGAs that we've been describing.

Surveillance UAV example

Battlefield surveillance UAVs such as the Altair Predator B variant (Figure 2) are good examples of where packaged PMC-based FPGAs like the TS-MPEG-4 could be used for video capture and compression. This class of UAV usually flies at low to medium altitude over a battlefield or other area of particular interest and carries a number of video and, possibly, high-resolution single-shot cameras for a more detailed view of individual objects. The UAV will be controlled from a ground station that receives images from various cameras and displays them for analysis by the ground crew. The images may then also be relayed further up the command chain to build a complete tactical

picture of the battlefield. The downlink from the UAV to the ground does not have the bandwidth to transmit all the video streams directly from the cameras in real time, driving the need for compression.

The mission computer for such a UAV is likely to be implemented using COTS VMEbus or CompactPCI modules. Because of the limited space, weight, and power budgets available in a UAV, 3U CompactPCI would again be an ideal format choice for the mission computer. FPGA-based PMC modules for video compression could be mounted on a host SBC or could occupy 3U slots using carrier cards. Video streams direct from the cameras in RS-170 format would be converted to MPEG-4 by the FPGAs, then encapsulated and downlinked by the mission computer for any of the ground-based operations required.

The FPGA with its unique and flexible architecture looks set to replace many of today's dedicated DSP solutions where its parallelism and aggregate throughput make possible big reductions in real estate and cost. Equally, the cost of time-to-deployment is becoming a critical factor for both the government and system inte-

grator, and FPGA-based solutions often provide benefits as well. The availability of bundled, application-oriented COTS solutions, even though they may require minor customization for a particular end-use, promise to bring new FPGA-based DSP systems online faster and at lower cost.

Duncan Young has worked in the defense industry for almost 40 years. Duncan was part of the management buyout team that formed Radstone Technology, and he initiated product development of conduction-cooled VMEbus modules. He has also served on a number of standardization committees. Duncan is now an independent consultant and writes this column on behalf of SBS Technologies.

For more information, contact:

SBS Technologies

2400 Louisiana Blvd. NE, Ste. 5-600

Albuquerque, NM 87110

Tel: 505-875-0600

Website: www.sbs.com



Figure 2

When failure is not an option.



0 mph



921 mph

Vibration issues?

We have the answer.

At 921 mph every part of the F-22 is put to the ultimate vibration test.

Every electrical connection has to work perfectly... the pilot is betting his life on it.



**Ruggedized
VME64x
Connectors**
Where reliability
and speed of
data transfer
are a necessity

All Hypertronics products are built upon the legendary Hypertac® Contact which outperforms other interconnect options.

- Immunity to shock and vibration.
- Up to 100,000 mating cycles.
- Nearly half the resistance of conventional contact designs.
- Extremely low insertion and extraction forces.


HYPERTAC
HYPERTRONICS
INTERCONNECT SOLUTIONS

ISO 9001, 13485, & 14001 certified

Hudson, MA

1-888-HYPERTAC

www.hypertronics.com

smiths

Visit Hypertronics at MEECC - May 16th & 17th - Long Beach, CA - www.mecc.com

RSC# 19 @www.mil-embedded.com/rsc



Why convert to a SAASM-based Global Positioning System?

By Ron Holm

In 1998, the Joint Chiefs of Staff selected the Selective Availability Anti-Spoofing Module (SAASM) as the security architecture to bring the Global Positioning System (GPS) to the next level, and issued the following mandate: As of October 1, 2006, all newly fielded Department of Defense (DoD) GPS systems will use SAASM-compliant Precise Positioning System (PPS) devices. Procurement of non-SAASM GPS user-equipment will be disallowed unless waived.

Despite a government mandate requiring all newly fielded DoD GPS systems to use SAASM-compliant PPS devices by October 2006, many military groups and other federal agencies continue to purchase receivers without SAASM compliance. Military users purchasing a GPS receiver without SAASM, or anyone who waits until the October 2006 deadline, is taking a security risk. Standard GPS service could be denied at any time via warfare tactics such as jamming or spoofing, and if this occurs, GPS receivers without SAASM will find it difficult to correct the situation quickly because the process to acquire SAASM-compliant receivers requires significant time for authorization and processing.

We are going to explain what SAASM is all about, why it's important to GPS receiver end-users, and why those who deploy non-SAASM receivers are putting their organization at risk right now – even though the deadline is less than a year away.

Sense of urgency

The need for improving GPS security came to the forefront even more this past December in an announcement by President George W. Bush in which he issued the Space-Based Positioning, Navigation, and Timing (PNT) policy. The PNT policy authorizes the improvement of the United States' capabilities to deny hostile use of any space-based positioning,

navigation, and timing services without unduly disrupting civil and commercial access. In the policy, the President specifically directed the Secretary of Defense to develop and maintain navigation warfare capabilities required to effectively utilize GPS services in the event of jamming or other interference by adversaries.

This announcement underscores the fact that the federal government is increasing the level of urgency to safeguard GPS. The pressure for government agencies and military units to convert to SAASM-compliant GPS receivers is bound to also increase dramatically. Along with the selection of SAASM by

the Joint Chiefs, the writing on the wall is clear: All defense agencies should begin converting to SAASM GPS receivers.

SAASM explained


To understand the risks and why it's important to deploy SAASM-compliant GPS receivers as soon as possible, a brief recap of GPS and SAASM will help. GPS has come to play a significant role in our everyday tasks during the past decade. GPS makes it possible to pinpoint the precise location of any person or place, anywhere in the world. It helps with everyday things such as how to drive our vehicles from Point A to Point B via onboard navigation systems.

Spoofing and encrypted coding

GPS spoofing involves the intentional sending of a fake GPS signal by a simulated satellite mimicking a legitimate GPS satellite. Spoofing produces a false reading in GPS SPS devices and, if properly executed, can introduce position and timing errors, disrupting navigation and communication systems.

The low-power GPS satellite transmitters deliver extremely low-strength signals (equivalent to 0.000000000000001 watt) to Earth-based GPS receivers that are vulnerable to jamming and spoofing. Jammers are inexpensive, unintelligent electronic devices that merely produce a higher-power blocking signal at the GPS frequency. Jamming is disruptive but usually detected by the GPS receiver as it stops tracking satellites. Spoofing requires more sophisticated, expensive equipment. Spoofing poses a particular security risk as it is often undetected by a GPS SPS device.

The key to preventing spoofing is to deploy a GPS receiver that can acquire encrypted GPS signals referred to as P(Y) coded signals, which are more robust and jam resistant. GPS satellites broadcast two signals: a civilian, unencrypted signal (referred to as C/A) that all GPS receivers can access, and the military encrypted coded signal P(Y). GPS devices in compliance with SAASM can receive and decrypt the P(Y) code (when keyed), which authenticates that the signal originated from the GPS satellites. The code frequency ranges are shown centered about L1 in Figure 1.



GPS has also become critically important to the military to identify the whereabouts of friends and foes, and it plays a crucial role in the success of military operations by providing precise time and frequency to communication systems. This allows military units to synchronize movements and ensure they communicate over secure frequency bandwidths that change on an irregular basis to avoid detection by the enemy.

Since GPS relies on low-powered frequency waves traveling from satellites to GPS receivers on the ground, the technology also lends itself to intentional jamming by enemies as well as unintentional or intentional jamming by allies. For example, the Civil Coarse Acquisition (C/A) code signal may be intentionally jammed by the US and other allies to allow only SAASM and legacy P(Y) receivers to access GPS. GPS is also susceptible to enemy spoofing – the deliberate attempt to mimic a legitimate signal and introduce erroneous position and time information.

To combat this situation, the US government launched a program in the 1990s referred to as SAASM. SAASM deploys anti-spoofing measures using cryptography to protect authorized users from false satellite signals generated by an enemy. To understand the reasons for SAASM, it helps to have an understanding of the components of the GPS system used by people, organizations, and governments throughout the world.

GPS mini-history

Through a satellite navigation system, GPS provides positioning and clock time to GPS receivers on the earth. Conceptualized in 1973, the first GPS satellite was launched in 1978, and in 1995 the system became fully operational. Today the system consists of 28 satellites orbiting 12,500 miles above the Earth.

GPS was originally intended as a military force enhancement system but now serves dual purposes: GPS has evolved to improve not only military security but also the accuracy of the position, velocity, and time of any object on earth – securely to military users and freely to civil users. GPS does this by offering two positioning services: Precise Positioning Service (PPS) for authorized military users only and Standard Positioning Service (SPS)

for everyone, including the military. SPS utilizes a simpler, unprotected C/A code that is openly available to commercial, civil, and military users.

The GPS signals are transmitted on two L-band frequencies: L1 (1575.42 MHz) and L2 (1227.60 MHz). The SPS service is provided on L1 and the PPS service on both L1 and L2.

Manufacturers of SAASM GPS receiver modules and the products that they are integrated into, referred to as *PPS Host Application Equipment (HAE)*, must work closely with the Key Data Processor Loading and Installation Facility (KLIF) under strict guidelines. After manufacturing the SAASM unit, the GPS receiver manufacturer ships the SAASM hardware to the KLIF for the loading of the Key Data Processor (KDP) crypto software.




Figure 1

The DoD relies upon GPS as the primary source for position, navigation, time, and time synchronization. Therefore, the GPS network was also built to allow for the deployment of security measures. Selective Availability (SA) is a security technique that involves the introduction of intentional errors into the GPS signal, which denies full system accuracy to SPS users. On 2 May 2000, however, the effects of SA were set to zero and it appears unlikely SA will ever be set higher. But the potential still exists, and this would degrade the accuracy of GPS for SPS users. Anti-Spoofing (A-S) utilizes cryptography and special keys to protect a GPS PPS receiver from receiving false satellite signals generated by an adversary.

The SAASM manufacturing and integration process

The detailed regulations that the US government has applied to the SAASM manufacturing process clearly demonstrate how serious GPS security has become and the need for immediate conversion to SAASM-compliant receivers.

After return of the SAASM device to the manufacturer, production test is completed and the unit is ready for sale to JPO approved customers.

Figure 1 shows the SAASM GPS receiver manufacturing process and integration into PPS HAE.

SAASM receivers support two key types to decrypt anti-spoofing and remove selective availability:

Physical-form red keys are classified, and distribution is closely protected since red keys are unencrypted.

Newer, black keys, on the other hand, are encrypted and unclassified. They can be distributed and loaded electronically, although paper tape distribution is still common. The decryption of the key only takes place within the secure SAASM module. Black keys may be renewable in the future via Over-The-Air-Rekeying (OTAR). Black keys make sense because they solve key distribution problems and are useless to

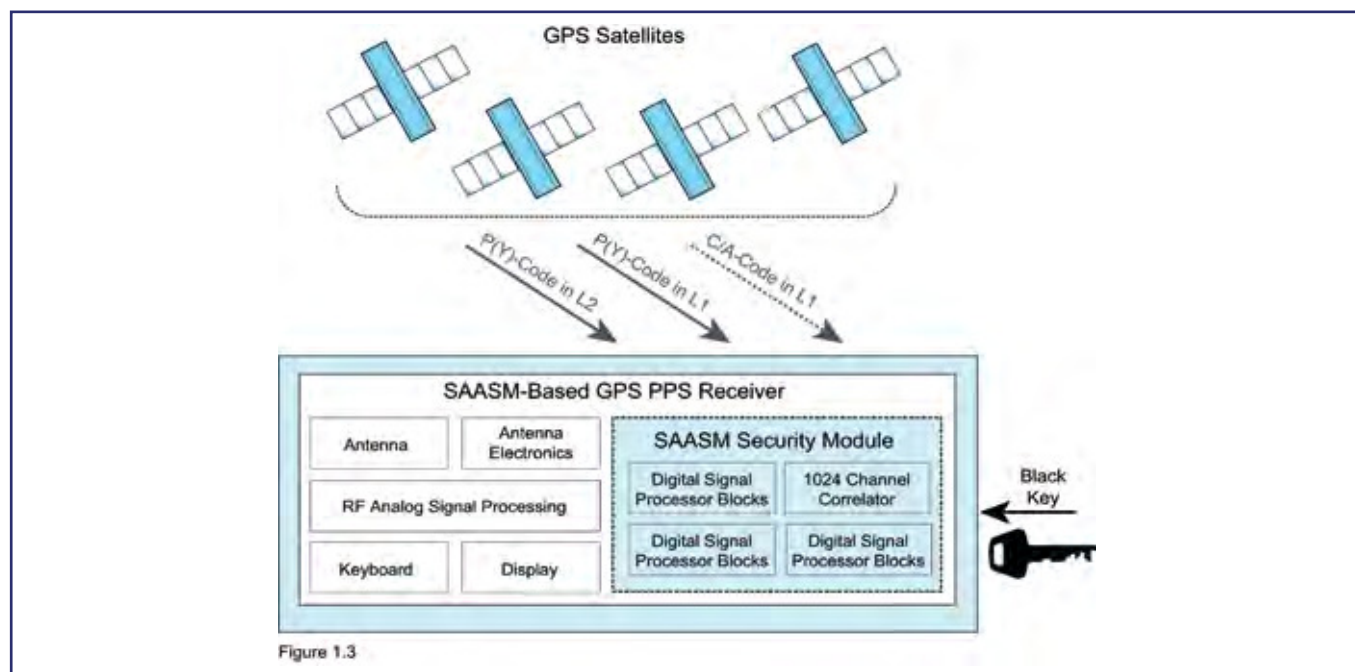


Figure 2

the enemy because they are encrypted. The DoD recognizes the security, delivery, and cost savings associated with black keys and wants users to transition from red to black keys as soon as possible (Figure 2).

To ensure security of the classified technology within a SAASM receiver and to preclude unauthorized procurement, the Navstar GPS Joint Program Office (JPO) at the Space and Missile Systems Center at Los Angeles AFB ensures compliance to DoD security requirements. All developers, integrators, and users of SAASM GPS must be JPO authorized.

SAASM developers and manufacturers must also meet several strict requirements, including securing a Communications Security (COMSEC) account, a KLIF account, and undergoing a complete JPO design-review process. Developers and manufacturers also have to prove they are free of foreign ownership, control, and influence and that they have a facility security clearance issued by the Defense Security Service (DSS).

Military end-users must receive authorization from the JPO to procure SAASM-based devices. When proper authorization occurs, the JPO issues a formal letter so that authorized manufacturers of SAASM GPS receivers will know authorization has formally been granted.

Benefits of SAASM GPS receivers

By purchasing a SAASM receiver now, authorized users comply with the DoD GPS security architecture and possess the most secure GPS technology to enhance their ability to use precise positioning, velocity, and time in all environments. SAASM-based GPS receivers have the capability to directly acquire the encrypted, military GPS code and no longer have to depend on the often-jammed civil GPS code.

The need for deploying a SAASM receiver could become critical at any moment since access to the Standard GPS service (SPS) could be denied via warfare tactics focused on local and regional denial (jamming) of the civil code within an area of conflict. In addition to losing access to the signal, upgrading GPS SPS based systems to SAASM is nontrivial since the process to acquire product and distribute the keys that can decrypt coded signals requires significant time for authorization and processing.

Without a doubt, those military and government agencies that wait until the SAASM deadline approaches in October 2006 are taking a security risk. If an enemy source attempts to jam or spoof the GPS signal, users of non-SAASM receivers could lose all of their GPS capabilities. †

Ron Holm is a product marketing manager for Symmetricom and manages the GPS Time and Frequency Instrumentation product line. Ron has been involved in the time and frequency industry since 1990 and has held previous positions in engineering and engineering management. Ron holds a bachelor's degree in electronic engineering from the University of Minnesota.



To learn more, contact Ron at:

Symmetricom

2300 Orchard Parkway

San Jose, CA, 95131

Tel: 408-964-7619

E-mail: ttm_sales@symmetricom.com

Website: www.symmetricom.com

The views expressed in this article are those of the author and do not necessarily reflect the official policy or position of the Navstar GPS Joint Program Office, the Air Force, the Department of Defense, or the U.S. government.

This article has been modified from the original October 2005 version to correct several technical inaccuracies. – Ed, July 2006.

CMSE 2006

Components for Military & Space Electronics
Conference & Exhibition



The CMSE is the first and largest International Conference on the use of electronic components in military and space systems. This annual Conference has the highest attendance of engineers, project personnel and offers the most technical papers on the following subjects.

Announcement and Last Call for Papers

Experts from throughout the world are requested to submit a one page abstract* describing the title, scope and content of the presentation while emphasizing the key points. Papers reporting on actual experience, knowledge and data gained from the use of commercial components and COTS in military and space systems are solicited. ***Specific subjects for consideration include:***

- Design Practices
- Applications
- Case Studies / History
- Obsolescence Management
- Radiation Hardness
- Reliability Analysis & Testing
- Embedded COTS
- Technology Trends
- Risk Mitigation
- EEE Parts & Components
- Selecting COTS and Commercial Suppliers

*Abstracts presented for consideration should be submitted before September 20, 2005.

CMSE 2006

February 6-9, 2006
Sheraton Gateway Hotel
Los Angeles, California

Organized by **Components Technology Institute, Inc.**

For more information on the Conference or Exhibition
please call 256-536-1304 or visit our web site at

www.cti-us.com

– Editor's Choice Products



Rackmount PPC 970FX-based HPC does 64-bit Linux

Increasingly, low-cost cluster computers running open source Linux are being used in semi-benign military environments. The Mercury XR9 1U server is an example of one of these systems. Mercury effectively "bolted" together two PowerPC 970FX-based high-performance computing cluster nodes for efficiency and speed. The XR9 runs Terra Soft's 64-bit Linux OS Y-HPC and cluster management suite called *Y-Imager*.

The Mercury XR9 contains twin 2.4 GHz CPUs paired up to an IBM CPC925 Northbridge with a 16-bit 400 MHz HyperTransport cave, providing peer-to-peer access between processors, memory, PCI-X bridge, or the Southbridge. There are two independent PCI-X buses, three PCI-X slots, and dual Gigabit Ethernet ports. There are also four SATA channels, four USB 1.1 ports, and dual RS-232 channels. There's an optional InfiniBand PCI-X HCA with dual 4X ports, and an optional serial FPDP PCI-X HCA add-in for dual 2.5 Gbaud ports. A separate PowerPC 405EP performs "service and control," freeing the main CPUs for number crunching.

Mercury Computer Systems

www.mc.com
RSC #23086

Terra Soft Solutions

www.terra softsolutions.com
RSC #23087

Hope for "classic" DS1691AJ/883 RS-422 drivers




When was the last time you saw a DIP-packaged IC? If you deal with mil systems, you might've been working with one of these babies as recently as yesterday. That's because 10+ year military life cycles require mature – and sometimes obsolete – devices to be sourced for a long, long time. We bring this to your attention because QP Semiconductor makes a business out of being a "Military Integrated Circuit Manufacturer" for QML and older SMD (DESC) devices. It may not be glamorous, but it's *essential* to our warfighters.

The company recently announced redesigned and replacement versions of the DS1691AJ/883 RS-422/423 line drivers, which originally came in both MIL-STD-883C and DESC 5962-8672101EA versions. Redesigned to match the original National Semiconductor 50V bipolar process devices, the replacement versions will be marked as QP1691AJ and are expected to be available as form, fit, and functional replacements for the originals. QP Semiconductor is performing a ground-up redesign; initial simulation models have been "successful," and the devices are in layout. Firm delivery dates will shortly be announced. The company maintains more than 1,540 ICs in their QML portfolio.

QP Semiconductor

www.qpsemi.com
RSC #23088



256-channel digital down-converter fits into Gateflow FPGA library

In order to process wideband radio frequencies in digital systems, Digital Downconverters (DDCs) are used once the RF has been converted at the system front end. Trouble is, the areas of interest – the *channels* – have historically required lots of DDCs because of low channel density per device. Pentek aims to change all that with their Model 4954-430 256-channel narrowband DDC intellectual property (IP) core. Designed for use with Xilinx Virtex II, Virtex II Pro, or Virtex-4 FPGAs, the company claims this to be "at least an order of magnitude" higher channel count than any other implementation available in FPGA or ASIC form.

The high channel-to-FPGA ratio can reduce cost, weight, and space on circuit boards and in systems. 16-bit real or complex data can be accepted at up to 185 MHz. The architecture utilizes a channelizer stage that generates 1K fixed adjacent frequency channels with alias-free signals of greater than 75 dB across the passband (for each channel). The core includes a 256-channel output switch matrix that gives the user the ability to conduct "coarse tuning" between channels. An NCO performs independent tuning on each channel. Factory default filter coefficients are available, or users can create their own. The Model 4954-430 fits into Pentek's existing and well-proven Gateflow FPGA design kit, which includes cores for data acquisition, FFTs, FIR, radar pulse compression, and other general and specific functions.

Pentek

www.pentek.com
RSC #23089



Palm-sized vehicle computer – small as an "ROC"

It was bound to happen – marrying the successful stacked concept of PC/104 modules with conduction cooling for harsher environments. Done

using conduction-cooled PMC and PrPMC modules by COTS supplier SBS Technologies, the company's Rugged Operation Computer (ROC) is designed for those tiny available shoehorn spaces found in every system. Weighing just six pounds and consuming only 100 in², the ROC is ruggedized for use in UAVs, ground vehicles, or even soldier/Marine portable use. SBS is targeting avionics, vetronics, and navtronics applications.

Available CPU modules include Intel or PowerPC processors on PrPMCs from SBS. I/O ranges from the company's broad line of PMC I/O, plus all industry-standard conduction-cooled PMC cards. The ROC accepts an integral stacked PMC structure, complete with a 100 W power supply, EMI filter, and a CompactFlash disk with up to 128 GB of memory mounted on a PMC carrier module. Operating systems range from Windows XP and Linux to INTEGRITY and VxWorks. This is one ROC you'll want to throw at your computing problem.

SBS Technologies

www.sbs.com
RSC #23090

Enabling military design security with high-performance FPGAs

By Joel Seely and Jie Feng



In an era of ever-increasing security concerns, SRAM-based FPGAs with a built-in configuration bitstream encryption feature offer critical advantages to military systems designers. In addition to high density, high performance, low development risk, and fast time-to-market benefits, they also deliver a secure approach for protecting designs against copying, reverse engineering, and tampering. For FPGAs without this built-in security feature, an additional nonvolatile device can be used to protect the FPGA design by supplying handshaking tokens.

Military applications are becoming increasingly complex. Major programs such as Future Combat Systems (FCS) and the Joint Tactical Radio System (JTRS) are pushing technological capabilities on all fronts to their limits. The electronics in these systems are relying on programmable logic and FPGAs to provide extreme flexibility at a reasonable cost while not giving up the requisite computational power. For example, secure communication systems are used to connect a variety of airborne, space, ground- and sea-based military communication networks. They are used in the transmission, processing, recording, monitoring, and dissemination functions of a variety of such networks, including secure data links. All this functionality requires processing power and reconfigurability.

SRAM-based FPGAs provide the highest density, performance, and most advanced features among all FPGAs, and they are reconfigurable. However, SRAM-based FPGAs are volatile and require external Flash memory or configuration devices to store their configuration data. At power-up, the configuration data is sent from the memory device to the FPGA, which can be intercepted. Copying, reverse engineering, and tempering are security concerns for these high-performance FPGAs used in military systems.

Two techniques – configuration bitstream encryption and handshaking tokens – can be used for securing designs within SRAM-based FPGAs. The configuration bitstream encryption is enabled using Advanced Encryption Standard (AES). Handshaking tokens are generated by a CPLD after communication with an FPGA. The FPGA only operates when the handshaking tokens from the CPLD match what it generates internally; otherwise, the design will shut down.

Let's investigate these techniques and how they secure SRAM-based military systems.

Configuration bitstream encryption

Some of the latest generations of SRAM-based FPGAs contain built-in AES decryption engines and key storage for configuration bitstream encryption. A generic flow of secure configuration can be carried out in the following three steps:

1. Load the AES decryption key into the FPGA.
2. Encrypt the configuration file with the same AES key and store it in the external memory, such as a configuration or Flash device.
3. At system power-up, the external memory sends the encrypted configuration file to the FPGA, which then uses the stored key to decrypt the configuration file in real time and configure itself.

AES comes in three different key sizes: 128-bit, 192-bit, and 256-bit. To understand the level of security, studies have shown that if one could build a machine that could discover a Data Encryption Standard (DES) key in seconds, it would take that same machine approximately 149 trillion years to discover a 128-bit AES key. (Source: National Institute of Standards and Technology)

Security key storage, which can be in either a volatile or nonvolatile location, is an important part of overall security. When the key is stored in volatile memory, an external backup battery is required when the power is down. Reliability, especially in military environments, is a major concern. Battery life depends on temperature and moisture levels of the surrounding area. If the battery dies, the key will be lost, and the device becomes unusable and must be sent back to the factory for repair. Batteries are often not desirable in military applications because of the possibility of leakage, outgassing, explosion, or going dead at an inopportune time. In addition, a battery increases overall system cost and requires special manufacturing attention. A nonvolatile key, which is more reliable, practical, and flexible, can be programmed during regular manufacturing flow with the FPGA either on- or off-board.

With configuration bitstream encryption, only encrypted configuration files and an AES key are stored in the system. Even if the configuration bitstream is captured, it cannot be decrypted without knowing the key. There are various ways to hide the AES key, including placing it under layers of metals in a distributed manner. Some FPGA vendors such as Altera also utilize additional security techniques to make the key difficult to find. Further, Altera does not support read back of the configuration file, whether it is encrypted or not. This adds another layer of security to the solution, making it very difficult to copy a design.

Other security breaches

Reverse engineering any FPGA design to a Register Transfer Level (RTL) or schematic format through configuration bitstream is very difficult and time-consuming, even without encryption. For high-performance FPGAs, the configuration file contains millions of bits. To reveal the mapping from the configuration

Hardware

file to the device resources, one needs to reverse engineer the FPGA or the FPGA development software. Some FPGA vendors do not disclose the configuration file formats, making reverse engineering more difficult. With configuration bitstream encryption, one needs to first find the key and decrypt the file. It may be easier and quicker to build a design from scratch than to reverse engineer a secured FPGA design.

Tampering is modifying the design stored in the device or replacing it with a different design. The tampered device may contain harmful design code capable of causing a system to malfunction or steal sensitive security data. Tampering cannot be prevented if a volatile key is used because the key is erasable; once the key is erased, the device can be configured with any configuration file. For the nonvolatile key solution, the device can be set to only accept configuration files encrypted with the stored key. A configuration failure signals possible tampering with the configuration file during transmission between the external memory and the FPGA, or during remotely communicated system upgrades. This is another advantage of a nonvolatile key.

Handshaking tokens

Configuration bitstream encryption is only available in high-density, high-performance SRAM-based FPGAs. The following solution allows any FPGA designs to remain secure against copying, even if the configuration bitstream is captured. This is accomplished by disabling the functionality of a user design within the FPGA until handshaking tokens are passed to the FPGA from a secure external device. Figure 1 compares the software license scheme with an FPGA security scheme.

Configuring the FPGA is similar to installing software onto a computer; the configuration bitstream is not protected. The external secure device is similar to the license file. The software will only operate when a valid license file is present. The user design within the FPGA will only operate when the handshaking tokens sent from the external secure device are valid. A simplified hardware implementation for this FPGA security solution is shown in Figure 2. In this example, a

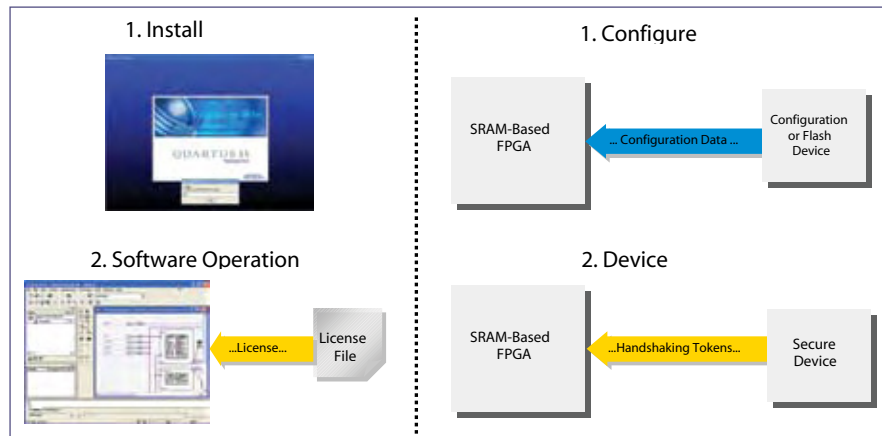


Figure 1

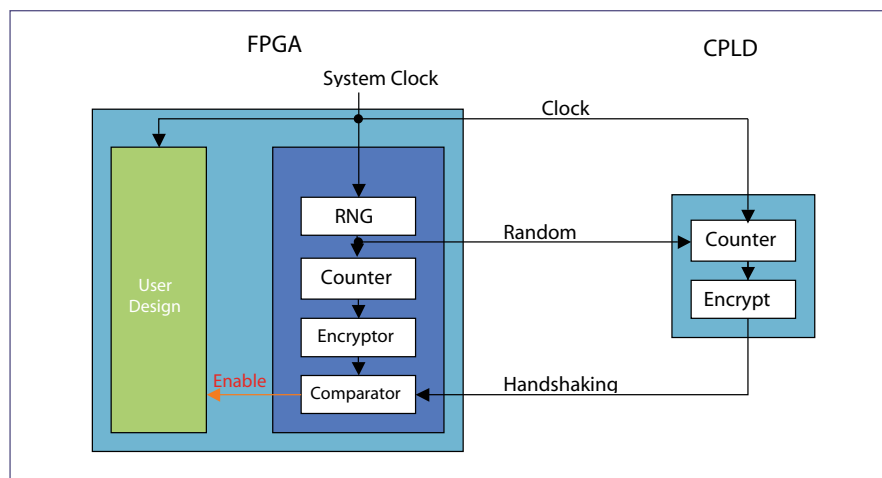


Figure 2

CPLD is used as the secure external device because it is nonvolatile and retains its configuration data during power-down.

After the FPGA is configured, the functionality of the user design within the FPGA is disabled because the enable signal is not asserted, while the security block within the FPGA starts to function. The Random Number Generator (RNG) generates and sends the initial counter value to the CPLD. The CPLD encrypts the counter value and sends the resulting handshaking token to the FPGA. If the handshaking token matches the data generated inside the FPGA, the enable signal is asserted and the user design starts functioning. This process continues during the entire operation of the FPGA. A mismatch will cause the enable signal to go low and disable the functionality of the user design. Figure 3 shows an example of how the enable signal is used with a simple AND gate.

The FPGA user design only works when the handshaking tokens from the external secure device and the data generated inside the FPGA are identical. Even if the FPGA configuration bitstream is stolen, it is useless, similar to software without a license. Therefore, the FPGA user design is secure from copying. This solution does not provide additional protection against reverse engineering (though difficult) and tampering.

The security of the solution relies on the external secure device being secure and the handshaking tokens being unpredictable. A secure external device needs to be nonvolatile and retain its configuration during power-down (for example, CPLDs or security processors). The RNG in the solution is critical. It ensures that every time the device starts up, it uses a different initial value. This prevents anyone from storing the handshaking tokens in a storage device. To prevent someone from

detecting the pattern in the handshaking tokens, a proven encryption algorithm such as AES should be used.

To ensure that the security scheme works properly, the system clock feeding the FPGA user design should be the same as the system clock feeding the security block. This prevents someone from disabling the security block when the enable signal is asserted. To further increase security, the comparator block can be duplicated several times to produce more enable signals to feed different portions of the user designs. †


Joel Seely has been with Altera Corporation for five years. Initially in charge of the embedded applications group, he is now technical marketing manager for the automotive, industrial, and military business unit.



For more information, contact Joel at:

Altera Corporation
101 Innovation Dr.
San Jose, CA 95134
Tel: 408-544-8122
Fax: 408-544-8066
E-mail: jseely@altera.com
Website: www.altera.com

Jie Feng is the product line manager in charge of the Stratix series FPGAs at Altera Corporation. She has been with Altera for more than six years and has experience in customer applications, wireless technical marketing, and product marketing. She has published and presented articles on FPGA design security at various conferences and in magazines. She holds a BS degree in computer engineering from the University of Toronto and an MS degree in computer science from Stanford University.



For more information, contact Jie at:

Altera Corporation
101 Innovation Dr.
San Jose, CA 95134
Tel: 408-544-6753
Fax: 408-544-6425
E-mail: jfeng@altera.com
Website: www.altera.com

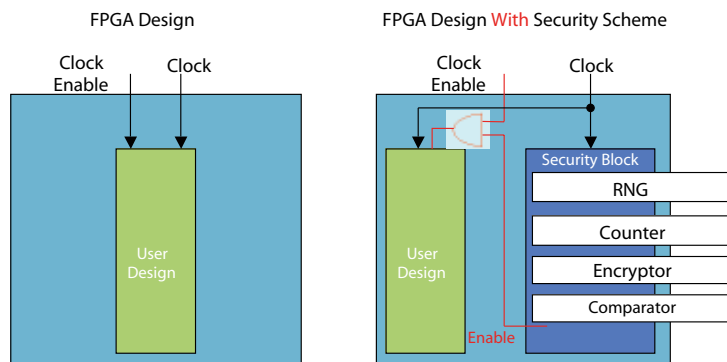


Figure 3

The most powerful DSP to ever hit the screen

BittWare's new 6U VME/VXS board. Ruggedized. Conduction-Cooled. COTS Compliant.

The military has specific needs. BittWare has the solution. The new T2-6U-VME board (T2V6), which supports the VITA 41 VXS (VME switched serial) standard, is the very first COTS VME-based board featuring eight ADSP-TS201 TigerSHARC DSPs from Analog Devices, providing up to 28.8 GFLOPs in a single slot. Designed and built for demanding real-time signal processing systems, the T2V6 is particularly effective in applications such as sonar, radar, imaging, and communications. Available in both air-cooled and conduction-cooled versions, it's ready for Airborne, Ground, or Naval environments. The T2V6 is the newest member of the T2 Family of boards from BittWare which includes PCI, cPCI, and PMC form-factors. Building on BittWare's decade-long expertise in SHARC-based products, the T2 Family is supported by a complete line of development software and run-time tools.

What's on your screen?

BittWare

www.bittware.com

TigerSHARC DSP is a registered trademark of Analog Devices, Inc.

RSC# 27 @www.mil-embedded.com/rsc



Software communication architecture: Evolution and status update

By Jeff Smith, Ph.D., David Murotake, Ph.D., and Antonio Martin

This article is a “must read” for anyone dealing with the SCA or its implementation. It gives a concise status overview, and it also identifies ties to other specifications and standards bodies such as OMG. - Ed.

Originally developed for the military's Joint Tactical Radio System (JTRS) program of Software Defined Radios (SDRs), the Software Communications Architecture (SCA) middleware is the key component to abstracting the underlying hardware from interoperable and programmable waveforms. The SCA, in effect, is what reprograms the radios to facilitate their reconfigurability. While the SCA remains strongly influenced by the JTRS program and the military, it's now also being considered for use in commercial and civilian applications. Throughout time, the SCA has evolved as users (civilian and military alike), industry committees, and complementary standards weigh in on SCA features and capabilities.

The SCA is a common specification standard and component-based software framework/architecture for SDR. The SCA is designed to facilitate waveform portability between different platforms and to leverage commercial standards, frameworks, and architectures to reduce development cost and improve reuse. Areas addressed by the SCA include waveform download, interoperability, operation and deployment on SDR devices, APIs (such as network layers, security, and common devices), and common component information.

While there are many interpretations of SDR, for the purposes of this article, external devices and the infrastructure composing the software bus will not be included as their future is independent of the SCA (for instance, they may be addressed by different standards bodies or deal with hardware migration). Instead, we will address the SCA by categorizing it into divisions of infrastructure, waveform support, services, device interfaces, heterogeneous processor, and security.

There are many permutations for a future SCA based on anticipated and existing commercial and government developments. To achieve future goals, it's key to address the challenges in future SCA development, commercialization, and adoption, and to summarize the current state of the SCA and future recommendations. Related commercialization and government standardization activities will certainly also affect the SCA efforts.

SCA evolution

SDR has a range of meanings today, depending on the types and number of hardware components that are replaced or upgraded by software. For simplicity, *SDR* is a term coined to describe a radio with a software-based physical layer that:

- Accepts fully programmable traffic and control information
- Supports a broad range of frequencies, air interfaces, and application software
- Supports over-the-air changes of initial configuration and waveforms

The SCA originated with the JTRS primarily to support SDR waveform portability for a new family of SDR tactical radios for the US military. The Software Defined Radio Forum (SDRF) assisted the JTRS Joint Program Office (JPO) in developing this open framework for SDRs, beginning with version 0.9 to the current version 3.0 (see jtrs.army.mil/sections/technicalinformation/fset_technical_sca.html) with its associated Application Program Interface (API), Specialized Hardware and Security Supplements. The Specialized Hardware Supplement is the main addition to SCA 3.0, which includes other improvements such as the elimination of reference counting and security supplement enhancements.

As the SDRF continues to support development of the SCA, it has sponsored the development of both an Open Source Reference Implementation (OSRI) for an SCA-compliant Core Framework (CF) as well as a compliant waveform based on FM3TR. The CF, based on a hybrid Java and C implementation, is available to SDRF members. An FM3TR waveform project is expected for completion later this year; in addition, the SDRF has developed requirements, use cases, Requests For Information (RFI), and Requests For Proposal (RFP). Typically, these technical products are voted and approved by the SDRF, then transferred via formal liaisons to other organizations such as the JTRS JPO and the Object Management Group (OMG), an object-oriented software standards organization.

For the last three years, SCA evolution has taken a parallel commercialization path in the Software Based Communications Domain Task Force (sbc.omg.org) within the OMG. In this forum, the domain-independent portions of the SCA, the bulk of the SCA, such as Lightweight Logging, Lightweight Services, Lightweight CORBA Component Model, Smart Antennae API, Digital Intermediate Frequency API, Deployment and Configuration of Components, and several Security Specifications, are in various phases of the standardization process as separate specifications. Development of these separate specifications allows commercial participation in related tooling and infrastructure.

Future SCA revisions should decrease in size and complexity as these OMG domain-independent specifications are completed and used as SCA references. This trend has already begun as the Lightweight Logging API was removed from the SCA, referencing the completed OMG version.

The OMG strives for SCA compatibility with its own software radio domain-specific version. Synchronization of OMG software radio specification improvements with the SCA has been achieved through liaisons and OMG member participation in the JTRS SCA Technical Advisory Group (TAG) revision process.

Existing SCA divisions

To simplify the categorization of changes, the SCA can be thought of in terms of current work and divisions as depicted in Figure 1. The current SCA 3.0 infrastructure manages the hardware radio components deployment by configuring devices and making sure they are ready, providing a standard store for configuration files, machine state, user attributes, and functional software, and offering a waveform structure, control, and binding framework for heterogeneous processors.

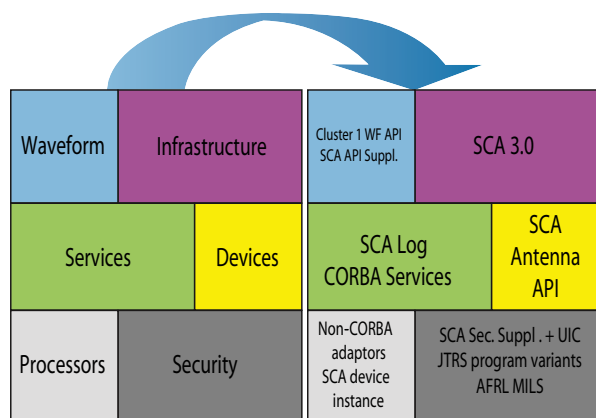


Figure 1

While not specifically addressing a waveform API, the SCA API supplement is given to support the portability of applications and interchangeability of devices; there is a specialization of the API derived from Cluster 1, a large SCA-compliant JTRS program. The current SCA assumes services that are provided by CORBA, for example, event and time services, and adds a logging service.

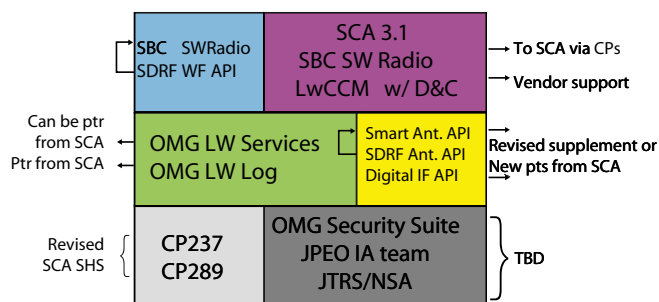


Figure 3

Device APIs, considered peripherals to the SDR, are also provided as an SCA supplement and at this time, an Antennae API is the only such supplement provided.

A standard method to access security functions such as encryption, authentication, transmission security (TRANSEC), and nonrepudiation, is specified in an SCA Security Supplement. Because of the nature of this technology, specializations exist for each JTRS program. In addition, there exists a parallel Air Force/NSA Multiple Independent Levels of Security (MILS) effort to combine the best of FAA DO-178B Common Criteria's security technologies, so as to provision secure services to embedded real-time, high-assurance platforms.

Parallel OMG standards plans and initiatives for the security functions and specifications are depicted in the OMG SCA Security Roadmap in Figure 2. A Specialized Hardware Specification SCA Supplement, available for SCA 3.0, specifies how to improve portability of software for processing elements other than general-purpose processors, including a Hardware Abstraction Layer (HAL) for deploying on heterogeneous processors.

Forecasted SCA divisions

Using the same divisions previously identified in Figures 1 and 2, Figure 3 shows a potential SCA evolution with possible choices. SCA changes occur through a Change Proposals (CPs) process and are reviewed through a Technical Advisory Group (TAG) and Change Control Board process. For instance, SCA 3.1 already

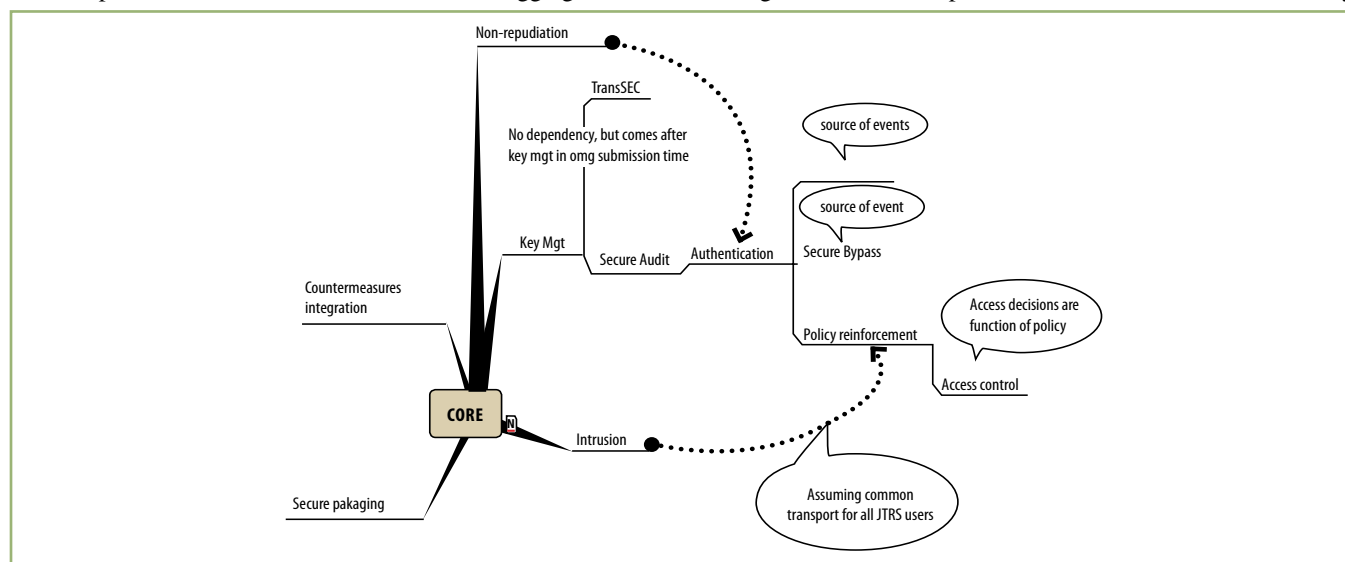


Figure 2

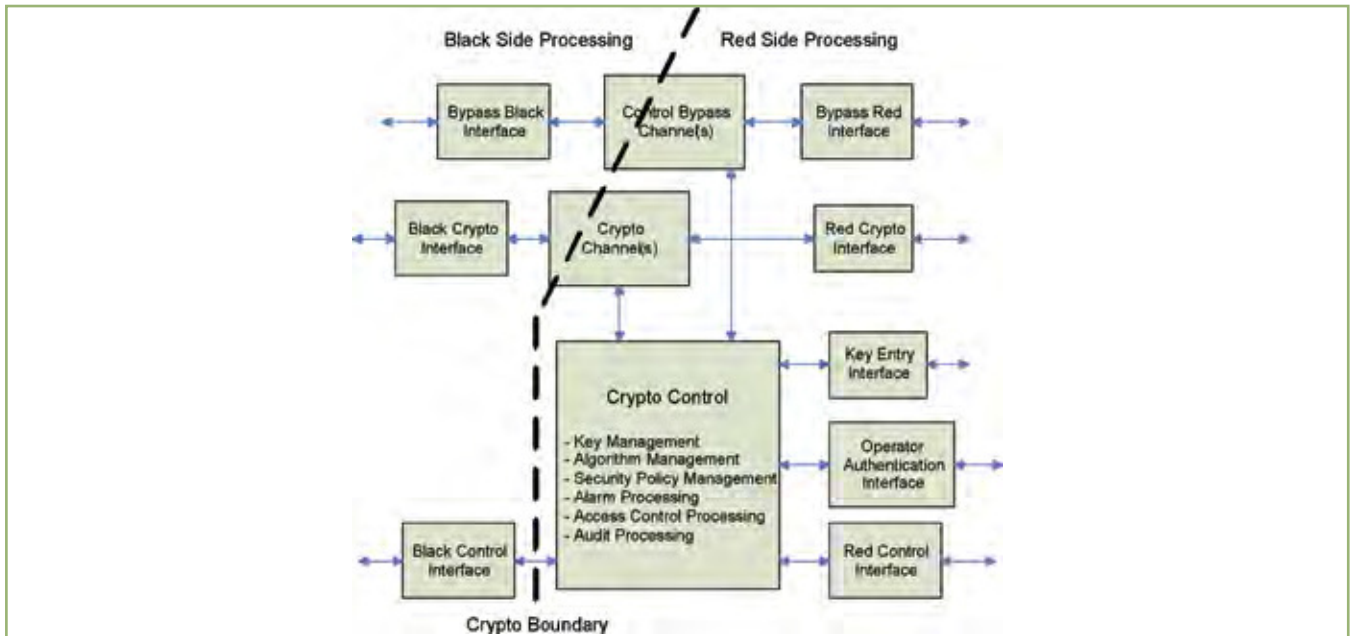


Figure 4

completed in draft form, includes CP289, detailing a Component Portability Specification (note CP289 was not accepted yet).

At this date, the OMG version of the Software Radio Specification is in the Finalization Task Force stage. This specification contains only the radio domain and waveform API portions of the specification, with the component model separated into different specifications that describe both the Deployment and Configuration of Components and Lightweight CORBA Component Model currently in the Revision Task Force stage. The SDRF is making additional progress with a new Waveform API contracted research and development project expected to begin in September, partially based on an OMG Software Radio Specification Waveform API Subset document. For the present, synchronization of the OMG and JTRS versions of the Software Radio specifications has been through OMG member participation in the JTRS CP process.

There are two new device-related specifications in process. The first is a Smart Antennae API Specification, with parallel efforts in both the SDRF and OMG. The second is an OMG Digital Intermediate Frequency (DIF) API Specification providing a standard API between tuners and the computer(s) hosting the rest of the software radio logic. This DIF specification is the software analog of the hardware standard driven through the digitalIF.org (www.digitalIF.org) standards group.

As previously mentioned, the OMG Lightweight Logging Specification has been finalized, serving as an SCA reference. The closest services specification to finalization is the OMG Lightweight Services Specification, offering a further reduction in SCA complexity.

The OMG Security Specification roadmap (refer back to Figure 2) is still in its initial phases; the first two specifications on this roadmap, the Core and Key Management Specifications, are in the initial submission stage. The OMG will standardize on the black, crypto, and red processing described in Figure 4.

Common security requirements are combined into this Security Subsystem Core to describe the overlap in one specification. The Secure Audit and Authentication RFPs are complete, with initial submissions in work; the rest of the OMG security submissions in Figure 2 will follow. In the meantime: 1) There are JTRS/NSA planned upgrades to the SCA Security Supplement; and 2) the Joint Program Executive Office (JPEO) is putting together an Information Assurance team to plan upcoming security specification update and implementation testing.

Tuning in the future

If the trend to replace SCA sections with domain-independent portions continues, tool vendor support will increase. In addition, the SCA framework will be smaller, require less testing, and eventually support ultra lightweight deployments in small and low-power consumption devices. The OMG SWRadio domain-specific specification will, in the short term, contribute to the SCA though the JTRS SCA change process. The progression from SCA 3.0 to SCA 3.1 will support true waveform component portability over heterogeneous processors.

Commercial SCA adoption is still hampered by many factors such as tool and predefined component availability. Future integration of SCA and commercial Software Defined basestation specifications holds promise as competing commercial standards and the SCA improve with liaisons and information sharing between groups.†



Jeff Smith received his Ph.D. from Northeastern University in computer systems engineering, his MS in engineering management from Southern Methodist University, his MS in computer science from East Texas State University, and his BS in computer systems engineering from the University of Massachusetts. He has a 30-year track record in the acquisition, management, research, and development of advanced development/technology programs and heads a consulting company, Composable Logic, supporting SCA Technica.

He is one of three co-chairs of the OMG Software-Based Communications Task Force and participates with the OMG Ontology working group. His primary expertise is in modeling/formal methods and applied complex systems (multi-sensor fusion or software radio applications).

For more information, contact Jeff at:

Composable Logic

PO Box 3148

Nashua, NH 03061-3148

Tel: 603-566-0124

Fax: 603-222-2098

E-mail: jesmith@ComposableLogic.com

Website: www.ComposableLogic.com



David Murotake received his SB in electrical engineering, SB in English literature and creative writing, SM in electrical engineering and computer science, and a Ph.D. in management of technological innovation from MIT. With more than 30 years of engineering and management experience at the US Army, RCA, GE, Lockheed, and Mercury Computer, he founded SCA Technica, Inc. in 2002. SCA Technica specializes in research and development of high-assurance SDR and CR and developed the High Assurance Wireless Computing System (HAWCS™) for protecting SDR and wireless computers from blended radio and Internet "hacking" attacks. A member of the SDRF board of directors, Murotake chairs its markets committee and is its Technical Committee's past vice-chair. He is founder and chair of numerous SDR working groups and special interest groups, including the Base Station WG and R & D WG.

For more information, contact David at:

SCA Technica, Inc.

PO Box 3148

Nashua, NH 03061-3148

Tel: 603-321-6536

Fax: 603-222-2098

E-mail: dmurotak@scatechnica.com

Website: www.scatechnica.com



Antonio Martin holds a bachelor's degree from the University of Maine in computer science and is pursuing a master's, certified by the Committee on National Security Systems, in computer science with a concentration in security at Boston University. Tony has eight years experience in the fields of biometrics, image processing, scalable distributed database and transaction processing, and network and system security. He is a participant at the OMG SW Radio FTF.

For more information, contact Antonio at:

SCA Technica, Inc.

PO Box 3148

Nashua, NH 03061-3148

Tel: 603-321-5220

Fax: 603-222-2098

E-mail: tony.martin@scatechnica.com

Website: www.scatechnica.com

Take your pick:

Pentium M® or PowerPC™

Rugged Embedded Computers for Land / Air / Sea Operations

- ECC RAM capabilities
- Fanless operation
- Low power consumption
- Rugged design
- Extended temperature option
- Long term availability
- PC/104-Plus expansions

These features make the MPL products an ideal and compact solution where high processing power, flexibility and high quality is required. MPL products are ready for military, aerospace, medical, oil/gas exploration and many industrial applications.

RSC# 31 @www.mil-embedded.com/rsc

info@mpl.ch  **www.mpl.ch**
High-Tech-Made in Switzerland

MPL AG, Täfernstr.20, CH-5405 Dättwil/Switzerland
Tel. +41 56 483 34 34, Fax +41 56 493 39 29
US Distributor: SIP Inc. Phone 480-513-8979

RTD Embedded Technologies, Inc.

"MIL Value for COTS prices"™



Intel Celeron cpuModules™



VIA Eden cpuModules™



8000 MIPS dspModules™

cpuModules™
-40 to +85°C

	Intel Celeron™				VIA Eden™				AMD Geode™				486DX	
	CME46786HX650HR	CME47786HX650HR	CMX46786HX650HR	CMX47786HX650HR	CME36786HX1000HR	CME37786HX1000HR	CML36786HX1000HR	CML37786HX1000HR	CMG26686HX333HR	CMG27686HX333HR	CME26686HX333HR	CME27686HX333HR	CMV6486DX100HR	CMV6486DX100HR
Bus	AT Expansion Bus	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	PCI Universal Expansion Bus		✓			✓		✓		✓		✓		
	PC/104-Plus PCI Bus Masters		4			4		4		4		4		
	APIC (add'l PCI interrupts)	9	9	9	9									
CPU and BIOS	CPU Max Clock Rate (MHz)	650	650	650	650	1G	1G	1G	1G	333	333	333	333	100
	ACPI 2.0 Power Mgmt	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Max Onboard DRAM (MB)	512	512	512	512	512	512	512	256	256	256	256	32	32
	RTD Enhanced Flash BIOS	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Nonvolatile Configuration	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Quick Boot Option Installed	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Fail Safe Boot ROM (KB)	768	768	768	768	768	768	768	768	768	768	768	384	384
	USB Boot	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Peripherals	Watchdog Timer & RTC	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	IDE and Floppy Controllers	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	SSD Sockets, 32 DIP	1	1	1	1	1	1	1	1	1	1	1	2	1
	Audio	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		
	TFT Panel TTL or LVDS			✓			✓		✓		✓	✓		
	SVGA Interface	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	AT Keyboard/Utility Port	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	PS/2 Mouse	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		
	USB Mouse/Keyboard	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		
	RS-232/422/485 Ports	2	2	2	2	2	2	2	2	2	2	2	2	2
I/O	USB Ports	2	2	2	2	2	2	2	2	2	2	2		
	10/100Base-T Ethernet	1	1	1	1	1	1	1			1	1		
	ECP Parallel Port	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	aDIO™ (Advanced Digital I/O)	18	18	18	18	18	18	18	18	18	18	18		
SW	multiPort™ (aDIO, ECP, FDC)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		
	ROM-DOS Installed	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	DOS, Windows, Linux	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

utilityModules™
-40 to +85°C

dspModules™

- Coprocessors
- Accelerators

Specialty I/O

- Pulse width modulator
- Incremental encoder
- Opto-isolated MOSFET
- User defined FPGA

Frame Grabbers

- Single or multi-channel
- MPEG-2 compression

Video Controllers

- Analog VGA
- TTL and DVI panel support

Communication Modules

- Copper or fiber Ethernet
- USB 2.0, Firewire, CAN, serial

Wireless Telematics

- GSM, GSM-R, CDMA
- EDGE, GPRS, SMS
- GPS, Wi-Fi, Bluetooth

Motion Controllers

- DC motor controllers
- Synchro, resolver, LVDT

Power Supplies

- 50/75/83/100 Watts
- Wide input range
- UPS backup
- MIL-STD-704/461

Mass Storage

- 1.8 and 2.5-inch IDE
- CompactFlash and PCMCIA



IDAN™—Intelligent Data Acquisition Node

- Build your own IDAN system
- Rugged PC/104 stacking concept
- Quick interchangeability and expansion
- Structural heat sinks and heat pipes
- Optional cooling fins
- Milled aluminum frames
- Optional MIL-SPEC paint
- Standard PC connectors
- Shock-mount isolation base



Full Product Line and Pricing Online

A Founder of the PC/104 Consortium • ISO9001:2001 Certified
Copyright © 2005 RTD Embedded Technologies, Inc. All rights reserved.

www.rtd.com

HighRel PC/PCI-104 Modules and Systems

-40 to +85°C



Autonomous SmartCal™



Wireless Telematics



Frame Grabbers

dataModules® -40 to +85°C		Smart A/D				Analog I/O				Digital I/O					
		SDM6440HR	SDM6540HR	SDM7540HR	SDM8540HR	DM6420HR	DM6430HR	DM7520HR	DM7530HR	DM6620HR	DM6812HR	DM6856HR	DM6888HR	DM6956HR	DM7820HR
Bus	AT Expansion Bus	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	PCI Expansion Bus Master			✓	✓			✓	✓						✓
	McBSP™ Serial Ports		✓	✓	✓			✓	✓						
Analog Input	Single-Ended Inputs	16	16	16	16	16	16	16	16						
	Differential Inputs	8	8	8	8	8	8	8	8						
	Max Throughput (kHz)	500	1250	1250	1250	500	100	1250	100						
	Max Resolution (bits)	12	12	12	12	12	16	12	16						
	Input Ranges/Gains	3/4	3/7	3/7	3/7	3/4	1/4	3/6	1/4						
	Autonomous SmartCal™	✓	✓	✓	✓										
	Data Marker Inputs	3	3	3	3	3		3							
Conversions	Channel-Gain Table	8k	8k	8k	8k	8k	8k	8k	8k						
	Scan/Burst/Multi-Burst	✓	✓	✓	✓	✓	✓	✓	✓						
	A/D FIFO Buffer	8k	8k	8k	8k	8k	8k	8k	8k						
	Sample Counter	✓	✓	✓	✓	✓	✓	✓	✓						
	DMA or PCI Bus Master	✓	✓	✓	✓	✓	✓	✓	✓	✓					✓
	SyncBus		✓	✓	✓			✓	✓						
Digital I/O	Total Digital I/O	16	16	16	16	16	16	16	16	16	48	32	64	32	48
	Bit Programmable I/O	8	8	8	8	8	8	8	8	8	24				48
	Advanced Interrupts	2	2	2	2	2	2	2	2	2	2				2
	Input FIFO Buffer	8k	8k	8k	8k	8k	8k	8k	8k						2M
	Opto-Isolated Inputs											16	48	16	
	Opto-Isolated Outputs											16	16		
	User Timer/Counters	2	3	3	3	2	2	3	2	3	3				10
	External Trigger	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				✓
	Relay Outputs													16	
Analog Out	Analog Outputs	2	2	2	2	2	2	2	2	4					
	Max Throughput (kHz)	200	200	200	200	200	100	200	100	200					
	Resolution (bits)	12	12	12	12	12	16	12	16	12					
	Output Ranges	3	4	4	4	3	1	4	1	4					
	D/A FIFO Buffer		8k	8k	8k			8k		8k					

RTD FieldPads™

- Ruggedized, embedded computer systems
- User-specified CPU and PC/PCI-104 expansion
- Weathertight components
- Integrated 6.5-inch video panel, keyboard
- Heat pipes for high performance CPUs
- User-defined MIL connectors
- Internal and external battery packs
- Expand with any RTD PC/PCI-104 product



Industrial FieldPad™

Ideal for control and monitoring of processes on factory floors or industrial installations. Mounting flanges allow the unit to be installed on machinery or walls, enabling standard PC access in a rugged enclosure resistant to industrial environments.



Tactical FieldPad™

Designed for mobile and portable applications where the angled panel and ergonomic design allow for optimal viewing with flexible positioning. Data collection/downloading and information access accomplished through wired or wireless connections.

HiDAN™ and HiDANplus™—HighRel Intelligent Data Acquisition Node

- HiDAN is a rugged, watertight enclosure for a stack of PC/104 modules
- HiDANplus combines the modularity of IDAN with the environmental ruggedness of HiDAN
- Integrated tongue and groove O-ring for environmental sealing and EMI suppression
- Structural heat sinks and heat pipes
- Optional cooling fins
- Milled aluminum frames
- Stackable signal raceway
- Optional MIL-SPEC paint
- MIL I/O connectors
- Shock-mount optional



www.rtd.com


Specifications, manuals, drivers, and plant tour

RTD Embedded Technologies, Inc.

103 Innovation Blvd • State College, PA 16803
T: 814-234-8087 • F: 814-234-5218

RSC# 33 @www.mil-embedded.com/rsc

rttd®
"Accessing the Analog World"®



Cognitive radios: The future of SDR technology

By Dr. Bruce Fette, Ph.D.

The future of two-way radio communication lies in the ability to use a single device to network with other types of devices while intuitively maximizing the power of flexible and adaptive software-based protocols. That future – the Cognitive Radio (CR) – is drawing closer, thanks to Software-Defined Radio (SDR) technology.

Cognitive radios will learn and autonomously perform “cognitive” functions as a form of intelligence that comes from their ability to be defined and upgraded using software. SDR is the foundation upon which cognitive radio will be built. To understand cognitive radio, we first must take a look at SDR.

The term *SDR* was originally coined by DARPA’s chief scientist, Dr. Joe Mitola, who saw a graduation of technologies that began with the hardware-defined radio and evolved into the digital radio and the software-defined radio in which all applications can be configured by software. The Software Defined Radio Forum (SDRF) has been working with this technology for several years and similarly defines SDR as a radio in which the software manages and controls the radio’s waveform properties and applications. Furthermore, an SDR is reprogrammable and may be upgraded in the field with new capabilities.

One of the SDRF’s functions is to define the standards by which those upgrades can be performed so that new technology can be harmoniously integrated into the radio after it has been fielded without completely replacing all the previous hardware functionality. These standards will allow equipment developers and, eventually, users to enhance the capabilities of their equipment. Because SDRs can be upgraded, bugs can be fixed and additional functions can be delivered to customers, creating incremental value.

SDR technology standardizes the architecture and supports a wide variety of modulation strategies, access strategies, and protocols as well as higher-level systems protocols such as trunked radio (a system used by local government and industry to operate private systems when a large number of mobile radios need to share frequencies), satellite communications systems, and even wireless access.

“[CR] represents an SDR with not only the ability to adapt to spectrum availability, protocols, and waveforms but the capability to learn waveforms and protocols, to adapt to local spectrum activity, and to learn the current needs of its user.”

The US military has embraced the properties and characteristics of SDR technology by way of the Joint Tactical Radio System (JTRS) initiative, which will bring communications interoperability to each of the armed services across all of their platforms. JTRS software-defined radios can fulfill multiple military communications functions and will ultimately be small enough to integrate into miniature robotic devices or body electronics worn by soldiers. Today, some 20,000 SDRs are being used by the US military and government agencies including the US Navy’s Digital Modular Radio (DMR), combat search and rescue radios, and specialized radios for government agencies in law enforcement and Homeland Security applications. SDR technology will improve interoperability among military services, coalition partners and public safety officers while solving bandwidth problems by reducing the number and types of radios required to accomplish operational objectives.

From SDR to CR

CR builds on SDR technology. It represents an SDR with not only the ability to adapt to spectrum availability, protocols, and waveforms but the capability to learn

waveforms and protocols, to adapt to local spectrum activity, and to learn the current needs of its user.

CR technology enables the radio itself to learn, allowing it to perform “cognitive” functions such as identifying and using empty spectrum to communicate more efficiently. CRs will sense and adapt their behavior according to the environment in which they operate. Once there is an embedded machine in which the software implements the protocols programmed for it, the radio is able to be smart and alert and it can “negotiate” with its environment. For example, a CR would learn about various services of interest to its user by being aware of its user’s activities. The radio knows how to find those services and knows the likelihood that some services will be of interest to the user in the immediate area. For instance, a CR could be aware of a Bluetooth network and what is available and of interest to its user within the Bluetooth service zone. It could also be aware of what’s available in a wireless LAN range, cell phone range and so on.

How does a CR get that smart? The defense community refers to a process called the OODA Loop – Observe, Orient, Decide, and Act. This is similar to the process humans perform as they go about deciding what to do in a situation. Those concepts can be extended to include planning and learning in the cognition cycle. The CR may do many of these kinds of things. It may observe and orient itself to the spectrum environment and decide and act on certain needs and wants of its user.

Academic research, industrial research, and research in the Department of Defense will synthesize new protocols, *etiquettes*, and technologies in the form of software that’s integrated into the CR. SDRs and cognitive radios must use etiquettes to know when it’s appropriate to interact and how to interact with their environment.



Take a closer look at your encryption.

Are you **getting the
protection you need?**

Members of the Sierra™ family of encryption modules are easily integrated into communications devices of all kinds. Sierra encrypts classified information, processes data at a higher speed, is extremely power efficient and allows modules to be reprogrammed as missions change.

Learn more at:

www.rfcomm.harris.com

www.harris.com

Broadcast
•
Microwave
•
RF
•
Government

HARRIS
assuredcommunications™

RSC# 35 @www.mil-embedded.com/rsc




Figure 1

Figure 1 shows a representation of what a cognitive radio knows:

- Where it is
- Which services are available – can identify and then use empty spectrum to communicate more efficiently
- Which services interest the user – and how to find them
- The current degree of needs and future likelihood of its user's needs
- How to learn and recognize usage patterns
- How to apply *Model Based Reasoning* about user needs, local content, and environmental context

Technology drivers and infrastructure

Key embedded technologies essential to evolving the SDR into a CR will include DSPs that, among other functions, manage modulation, cryptography, protocols, and source coding for voice, data, and imagery. High-density FPGAs are enabling shared, in-system reconfiguration and are the workhorses that change waveforms and adjust performance characteristics, frequency, power regulation, and other attributes. General Purpose Processors (GPPs) must manage more complex modem and operating environ-

ment controls that include the Software Communications Architecture (SCA), CORBA, and Real-Time Operating System (RTOS). GPPs will need to have significantly increased processing functionality while keeping size and power consumption to a minimum.

As the embedded enabling technologies advance core radio functionality, the next step in CR will be the application of two basic communication “etiquettes”: infrastructure and spectrum awareness. These two etiquettes will determine *when* it's appropriate to interact and *how* to interact with the communication environment – the “smart” part of the cognitive radio.

Infrastructure supports the radio's ability to manage policy, which is the regulatory governance (like the FCC) that defines user requirements pertaining to which frequencies are available for which purposes, which power levels may be used, and modulation and access permissions. Historically, infrastructure has proven to be a powerful tool in improving communication

system performance, particularly for trunked radios, cellular spectrum borrowing, and demand-defined multiple access. For the defense community, the most common infrastructure is Demand Assigned Multiple Access (DAMA). More than 30 years old and fully mature, the DAMA system processes and grants military communication user requests for a certain amount of time and bandwidth, thereby enabling a voice, data, image, or other communications events.

Spectrum awareness is based on waveform *orthogonality*, meaning that waveforms are intentionally designed to minimize interference between multiple users. Waveform differentiation can be found in time, frequency, code, modulation, and antenna beam-forming techniques. Orthogonality is particularly important for military applications, such as the Navy's DMR.

In Figure 2, the DMR can identify a number of existing signals in the frequency, spectrum, space. It identifies the frequency and modulation type of the signal and it tracks the signal. At the top of the screen in the figure, there is a green and red bar. The activity levels of the signals from the last 10 seconds are measured, providing a way to identify an occupied or unoccupied channel. When the user releases the channel, it is immediately free for another user on the same frequency. It also indicates that new waveforms may occupy the open frequency in this spectrum space. Orthogonality is particularly important considering the number of waveforms associated with military and coalition communications.

What CR will do

To recap, CR infrastructure awareness indicates the radio's ability to operate



Figure 2

according to policy. Spectrum efficiency optimizes radio performance by a number of factors based on orthogonality in the dimensions of time, frequency, code, or modulation.

The cognitive radio will also be capable of sensing, responding, and determining optimal responses to network and geographic operating conditions. Take, for example, an ad hoc network, which is a popular and pervasive network architecture (Figure 3). For an ad hoc network, when node "A" wishes to communicate with node "Z," it doesn't need to generate a *transmit* signal strong enough to cover the entire distance between the two nodes. Rather, node "A" sends a signal strong enough to communicate to an available node that is along the path to "Z"; this reduces power and saves energy. By moving along the path from open node to open node, the information bypasses potential delays caused by the transmission "waiting" for "busy" nodes to become available.




Figure 3

In addition, cognitive radios will be aware of subtle nuances within the network's structure such as the physical environment that includes data links, transport, and management layers. Figure 4 provides an example of a protocol stack and illustrates the interoperation and bridging of existing defense communication networks.

Geographic awareness is significant, particularly for international and coalition communications. This awareness generates the radio's ability to "discern" local infrastructure or policy, transmitters and receivers, terrain, propagation channels, and the location of network members.

FTP, TFTP		TELNET		SMTP			TNS
BGP		SNMP		MIB Extensstions		BOOTP	
TCP							UDP
IPSO, ICMP, IGMP, SDB, Tailored SNDCE							
Mil-STD-188-220		X.25/x.75 inc ADDSI			Mil STD 110A	Mil STD 141A	IP-SDR
Mil-STD-188-220		HDLC-LAPB/LAPD			Mil STD 110A	Mil STD 141A	HDLC LAPB/LAPD
RENAD Layer	HQ-II-NADLayer	EPLRS NAD Layer	Non-DAMA Access Layer	Non-DAMA Access Layer	Mil STD 110A	Mil STD 141A	WNW MAC
Sincgars SIP	Havequick II	EPLRS VHSIC	WB TACSAT	SATCOM DAMA	HF Modem	HF ALE	WNW Like

Figure 4

As an example, a US Air Force jet is flying across European airspace. Each country has different communication standards, frequencies, and protocols. With a cognitive radio, the flight plan is programmed into the radio, and just as GPS tells the pilot where they are, the radio would adopt the communication architectures of the airspace throughout the flight – without pilot interface.

"The cognitive radio will also be capable of sensing, responding, and determining optimal responses to network and geographic operating conditions."

Another link in the chain of awareness for CR is its ability to understand a user's operating requirements. Through software, the cognitive radio could reference the user's rank, role, and access requirements along with the databases and networks needed to complete normal operating tasks, thus keeping information consistent with the user's mission.

Additional cognitive abilities would authenticate and certify system access. Defined and managed by the user or the user's superiors, it can be accomplished over the air. Next, speech and language identification can be added. This is a component of functional awareness where the radio understands the syntactic and semantic context of dialog and can switch back and forth between text and speech, even performing biometrics on the user.

What's next?

Supporting and maintaining the communications security policy is also critical for the protection of personnel and data; it is an integral part of the radio's cognitive functions. Understanding and implementing national security policy, network operator, hardware/software, server authentication policies, stability and performance assessments round out the majority of a CR's functions.

The work of the SDRF, with its associated workgroups and committees, will bring the technologies and associated policies necessary for CR to standardization and reality. +

Bruce Fette, Ph.D., is chief scientist for the communication networks business area of General Dynamics C4 Systems in Scottsdale, AZ, and has worked for 34 years in advanced signal processing for telephony and RF communications. Fette currently serves as the SDRF's technical chair and is a panelist for the IEEE Conference on Acoustics Speech and Signal Processing Industrial Technology Track. He holds 35 patents and has been awarded the Distinguished Innovator Award. He received his BS in electrical engineering from the University of Cincinnati in 1969, his MS in electrical engineering from Arizona State University in 1974, and his Ph.D. from Arizona State University in 1981.



For more information, visit the General Dynamics C4 Systems website at www.gdc4s.com.



Mapping waveforms to systems: What would a wideband networking waveform system require?

By Kevin Maier

When building a Joint Tactical Radio System (JTRS) modem that must support a wide variety of disparate waveforms, it can be very challenging to assess the overall system requirements. A pragmatic process exists that can be employed to evaluate the requirements of a wideband waveform and to map these requirements to a hardware platform.

The steps recommended in this process include:

1. *Creating a representative waveform with the most demanding attributes*
2. *Performing a functional analysis of the target radio architecture*
3. *Mapping the functional analysis to specific processors*
4. *Determining the data flow, bandwidth, and latency requirements*
5. *Assessing requirements against hardware platforms*

The data presented herein is intended to be instructive in nature and is not exhaustive.

From a functional perspective, most modern tactical military waveforms can be classified into three broad representative waveform classes that all emphasize different parameters, and a rigorous platform evaluation can be done for a representative waveform in each class:

- Narrowband slow hopping waveforms (< 4k hops per second and <1 Mbps data rate)
- Narrowband fast hopping waveforms (> 4k hops per second and <1Mbps data rate)
- Wideband networking waveforms

While wideband networking-type waveforms are typically the most resource-intensive and demanding of all three classes, in reality, the first two classes should also be assessed if the JTRS modem is to support them. However, the following platform evaluation and information is limited to the examination of a representative waveform in the wideband networking class only.

Step 1: Creating a representative waveform

A wideband networking waveform may have a broad variety of parameters. For instance, the Joint Tactical Radio System Wideband Networking Waveform (WNW) has four distinct signals in space to fulfill its operational requirements: Bandwidth Efficient Advanced Modulation, Orthogonal Frequency Division Multiplexing, Anti-jam, and Low Probability of Intercept. For this illustration, we will create a waveform representative of the most demanding parts of the WNW subcategories.

Our representative waveform would enable end users to exchange wideband data on the battlefield or, with systems providing reach-back capability such as the WIN-T program, through the Global Information Grid (GIG). Other examples of this waveform type include the Tactical Data Rate waveform supported by the TRC-4000 terminal from Thales, and the AN/GRC-245 HCLOS terminal from Ultra Electronics[1].

Many of the parameters associated with JTRS and other waveforms are not publicly available; however, many of these waveforms are similar to the IEEE 802 series of networking waveforms (802.11, 802.16). As such, we will model the waveform based on publicly available information and supplement that information with details from 802.x waveforms. This hybrid structure will not specifically meet with the requirements of either 802.11 or WNW, but will be sufficient to drive the requirements of our modem architecture and help to evaluate against a hardware platform. The parameters for this representative waveform are presented in Table 1.

Parameter	Value
Channel Location in Spectrum	225-2000 MHz
Channel Spacing	30 MHz
Modulation	OFDM, DQPSK
INFOSEC	N/A
TRANSEC	None
Hops Per Second	None
OFDM Symbol Length	80 Samples
OFDM Guard Interval	16 Samples
OFDM FFT Window Length	64 Samples
Baseband Sample Rate	30 Msps
Error Correction for Payload	Rate 1/2 Turbo Code
Uncoded data bits per OFDM symbol	48
Signaling (bps)	18 M
Packet Structure	10 Training Sequence Bursts, 1 long OFDM Synchronization Symbol, 1 SIGNAL Header Symbol, 1 to N Data Symbol
Access Structure	Carrier Sense Multiple Access
Acknowledgement Time	10 μ s
Duplex	Full

Table 1

Our representative waveform is connection-oriented, meaning that for each packet transmitted, an acknowledgement must be sent indicating whether the packet was successfully received or not. This acknowledgement is based on a test of the Cyclic Redundancy Check (CRC) embedded in the data. The CRC response must be transmitted 10 μ s following the receipt of the last OFDM symbol, similar to the acknowledgement time for 802.11g. A hard requirement such as acknowledgement time will help drive the latency requirements of our target system.

Step 2: Performing a functional analysis

After creating a representative waveform, the next step in the platform evaluation exercise is to construct a functional block description of the target radio architecture. The wideband representative waveform can be deployed in a system similar to Figure 1, which shows the black side (unclassified) of the modem that partitions the modem functionality into six arbitrary blocks:

- RF front end – Provides the antenna interface and first stage of down conversion, accepts and distributes IF data to the next functional block
- IF processing – Performs analog-to-digital conversion and up/down conversion of the digital data
- Network synchronization and control processing – Converts data to/from symbols, contains the Finite State Machine (FSM) that controls the radio
- Receive channel processing – Performs demodulation, deinterleaving and deframing, and FEC decoding

– Provides higher layer data protocol processing and interfaces with the red-side modem processing

- Transmit channel processing – Performs modulation, interleaving and framing of data, as well as FEC encoding

In a supporting analysis, the exact data rate, latency, signaling, and pin count (where appropriate) are determined for particular data flows (indicated by letters) as appropriate to the representative waveform. A partial analysis of the receive path is shown in Table 2.

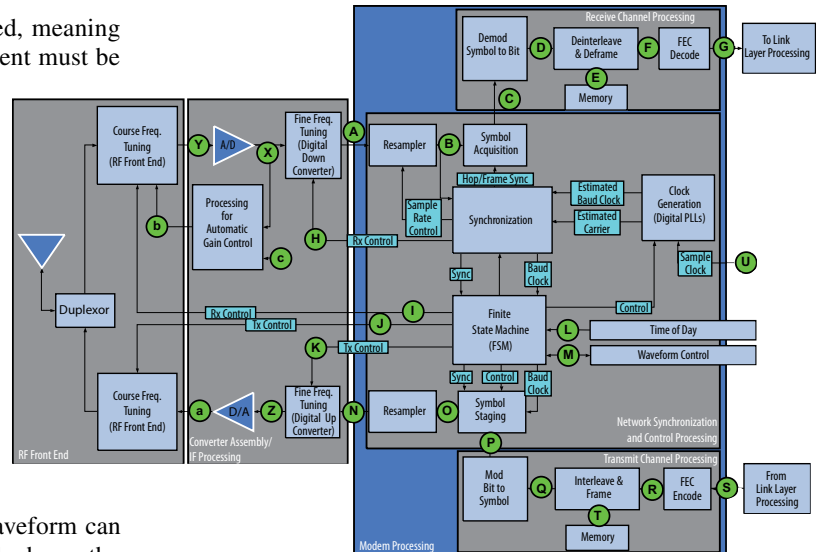


Figure 1

Path ID	Path Description	Bandwidth, Latency & Determinism Notes
A	DDC to Resampler	The baseband sample rate is 30 Msps, x 2 to allow for synchronization, x 4 bytes per sample (I & Q), which gives a rate of 240 MBps into the resampler
B	Resampler to Symbol Acquisition & Synchronization	The resampler will take the data rate down by a factor of 2, to 120 MBps
C	Symbol Acquisition to Demod	Assumption: the training sequence and frequency offset estimation symbols are generally irrelevant in the calculation of these rates. 30 Mbaud x 48/80 (packet overhead factor) x 4 bytes per sample (2 bytes per complex sample) = 72 MBps
D	Demod to Deinterleave and Deframe	2 samples per bit and 2 bytes per sample = 36 MBps
E	Deinterleave to Memory	2 x data rate = 72 MBps assuming 8-bit addressable regions. Could be up to 4x this rate if we only have 32 bit addressable regions
F	Deinterleave to FEC	36 MBps
G	FEC to Link Layer Processing	Rate ½ encoding gives 18 MBps
H	DDC RX Control from FSM	The fine control requires adjustment inside of the guard band of a symbol – or 20% of a symbol length = 0.533 μ Sec

Table 2

Step 3: Mapping to specific processors

After our representative waveform has been created and mapped to a functional block diagram of the modem, the results need to be allocated to actual physical processing elements. Notice at this point a target platform has still not been identified. The functional blocks should be assigned to specific processor types to estimate resource utilization that will help determine part sizing, leading to appropriate hardware platform architectures.

In general, high-speed or computationally intensive algorithms are mapped to FPGAs, with other algorithms mapped to either a DSP or a General Purpose Processor (GPP), depending on the anticipated power utilization and requirements for code portability. Higher layer back-end processing (link layer, network layer) is typically performed exclusively on a GPP.

In the case of this representative waveform, the functional processing blocks described above are mapped into all three different categories: DSPs, FPGAs, and GPPs. The functional blocks are partitioned as shown in Figure 2.

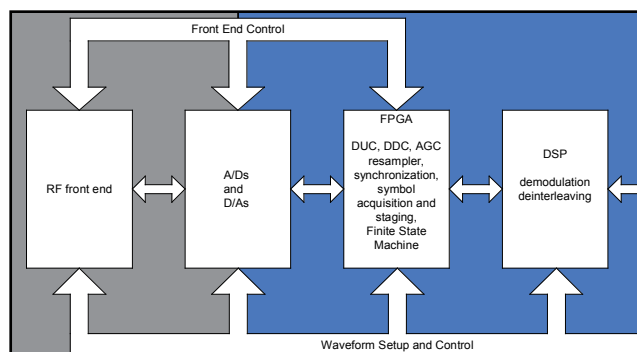


Figure 2

Once the waveform algorithms are partitioned across different processors, it is time to start our resource utilization assessment. Utilization estimates vary, but significant information is available in the public domain (for example, the Xilinx website provides FPGA utilization estimates). Much information about MIPS processing requirements of waveform components is also publicly available from the US government.

A detailed analysis of each functional processing element must be performed to create an accurate resource utilization estimate. For instance, our detailed receive side resampler analysis might be described as follows:

Rx resampler: Resampling is required to ensure that 4 samples per baud are input to the symbol acquisition and synchronization blocks. The resampling architecture follows that of the GC3011A resampler. The input rate to this resampler will nominally be 60 Msps, and the output rate will be 30 Msps. It is anticipated that the interpolation filter will be implemented using a Farrow filter, following the method prescribed by Harris and Dick[2]. This implements a 256-tap interpolation filter as five 8-tap filter stages with four MAC stages, thus the equivalent performance for the GC3011 is estimated to require sixteen 15-tap filter stages and 15 MAC stages. The 15-tap FIR filters can be implemented using a Distributed Arithmetic (DA) approach with a total estimated resource utilization of less than $16 \times 125 \times 2.25 = 4500$ slices[3].

Total estimated resource utilization would include 15 multiply and add blocks, and 1 block RAM in a Xilinx Virtex-II or Virtex-4 device. (Note that in Virtex-II, the adder would have to be implemented in gates). Also note that since the original filter architecture was based on XC4000 series FPGA technology, additional savings could be achieved using Virtex-II or Virtex-4 devices. As such, adding an additional 20 percent logic for overhead glue logic and allowing for floor planning risk puts total FPGA resource utilization at less than 5400 logic cells, 15 multiply and accumulate blocks, and 1 block RAM. Other detailed analysis leads to the summary presented in Table 3.

The detailed breakdown in Table 3 leads us to target a platform that has a single 600 MHz TMS320C6416 DSP, a GPP capable of at least 1000 MIPS, and an FPGA with approximately 32k logic cells and 3 x 18k block RAMs. In a size-, weight-, and power-constrained environment such as tactical MILCOM, the target platforms should encompass devices that do not have much excess capacity beyond a planned risk margin. For instance, one would not employ a Motorola AltiVec engine to field this waveform because of the high power consumption and general overcapacity of that particular processor.

Component Name	Proposed Processing Technology	Estimated Resource Utilization
Rx Resampler	FPGA	5400 Logic Cells, 15 MAC Blocks, 1 Block RAM
OFDM Symbol Synchronization	FPGA	3197 Logic Cells
OFDM Symbol Acquisition	FPGA	1590 Logic Cells, 2 MAC Blocks, 1 Block RAM
OFDM Symbol to Bit	FPGA	10744 Logic Cells
FSM	FPGA	5580 Logic Cells
Deinterleave/ Decode/ Interleave/ Encode	TMS320C6416	<100%
Link Network Layer Processing	GPP	< 1000 MIPS
Symbol Staging	FPGA	Minimal
Tx Resampler	FPGA	5400 Logic Cells, 15 MAC Blocks, 1 Block RAM

Table 3

There exist several modern processors that can fulfill the GPP requirements for target MIPS and target power consumption; for example, 1000 MIPS would require ~80 percent of a MPC 8541 at 533 MHz and run at approximately four watts. To fulfill the FPGA requirements, we could look at the two most recent Xilinx FPGA architectures; our waveform would require approximately 70 percent of a Xilinx Virtex-II Pro V2P40 or 53 percent of a Virtex-4 LX60. The Virtex-4 runs at approximately 50 percent of the power of a Virtex-II Pro and would be a suitable choice for our modem platform.

Step 4: Determining data flow, bandwidth, and latency

Once the functional blocks have been allocated to individual devices, the data flow requirements between our target devices can be summarized as shown in Table 4. This analysis will help determine if an interprocessor link requires channelized data flow and also what amount of bandwidth and latency is required for the link. The bracketed letters refer to the data flows identified in Figures 1 and 2.

Step 5: Assessing requirements against hardware platforms

Performing all of the above stages of analysis will generate a set

of hard requirements for a target platform. One can then choose to build the hardware from scratch or utilize existing Commercial Off-The-Shelf (COTS) boards and modules. The latter option is usually desirable if there are schedule or resource limitations. A first and easy test to see if specific COTS equipment is suitable is to use a checklist approach. Gather product information (typically available on a datasheet) and compare how well the product meets the overall system and data flow requirements. (Note that COTS designers marketing tactical MILCOM targeted systems should perform the same analysis to ensure that their systems can handle the target waveform requirements.)

The final step in assessing any platform is to actually build a representative waveform and port it to the target hardware. An example is the recent implementation of the 802.11g waveform on Spectrum's off-the-shelf SDR-3000 platform. This waveform requires the same bandwidth and latency performance as military wideband waveforms and as such demonstrates the SDR-3000's suitability for wideband waveform implementations.

The process described herein was employed during the development of Spectrum Signal Processing's newest series of Software Defined Radio (SDR) products. These SDR products have been designed from the ground up to support tactical

MILCOM applications; consequently, it was critical to validate that the modem architecture could support all of the JTRS waveforms in all of the key parameters: processor types and capacity, interprocessor bandwidth and latency, and external connectivity. +

References:

1. Pucker, L. "Paving Paths to Software Defined Radio," *Communication Systems Design Magazine*, June 2001.
2. Dick, Chris and Harris, Fred, "FPGA Interpolators Using Polynomial Filters," The 8th International Conference on Signal Processing Applications and Technology, Toronto, Canada, Sept. 13-16, 1998.
3. Xilinx LogiCore Distributed Arithmetic FIR Filter, v9.0, Product Datasheet DS240 May 21, 2004.

Kevin Maier has been a systems engineer at Spectrum Signal Processing Inc. for eight years, working on hardware, software, and embedded systems. He earned a BAS (honors) degree in systems engineering from Simon Fraser University.



To learn more, contact:
Spectrum Signal Processing Inc.
 #300-2700 Production Way
 Burnaby, BC V5A 4X1
 Canada
 E-mail: sales@spectrumsignal.com
 Website: www.spectrumsignal.com

Path	Channels	Aggregate Bandwidth	Aggregate Latency
FPGA to RF	Data: None	N/A	N/A
	Control: (I,J)	Low (under 100 kB/s)	All paths must be deterministic and bounded (~1-10 millisecond is appropriate).
RF to FPGA	None	N/A	N/A
FPGA to IF	Data: (A)	240 MBps	.5 μ s
	Control: None	N/A	N/A
IF to FPGA	Data: (N)	240 MBps	.5 μ s
	Control: (H K), U	Low	All paths must be deterministic. Channels H and K must be fixed latency
FPGA to DSP	(D)	36 Mbps	Low
DSP to FPGA	(Q)	36 Mbps	Low
FPGA to GPP	(M)	Low	Low
GPP to FPGA	(M)	Low	Low
DSP to GPP	Data: (G)	18 Mbps	Low
	Control: none	N/A	N/A
GPP to DSP	Data: (S)	18 Mbps	Low
	Control: none	N/A	N/A

Table 4



FPGA memory controllers improve DSP performance

By Richard M. Mathews

Signal processors spend a significant portion of time and resources moving data, shuffling it in preparation for manipulation. This inefficiency can be significantly reduced for downstream DSP processors by using a large, multi-ported memory buffer tightly integrated with a user-programmable FPGA logic block and a corner-turning Direct Memory Access (DMA) engine. This allows DSP and other processors to spend a higher percentage of time and resources on intelligent data manipulation, reducing overhead and system complexity. This article examines design issues and technology advances that can increase efficiency and optimize performance.

FPGAs are being increasingly used in DSP applications. They are especially effective at performing many kinds of repetitive operations and typically are combined with general-purpose processors for control operations and for data processing operations that require more complex decision making.

A common I/O path bottleneck problem exists in many DSP systems when moving data from node to node. Another problem occurs when a significant fraction of processing power is utilized for merely shuffling data in preparation for processing instead of using CPU resources for actual processing.

These data movement problems can be significantly alleviated by using a multi-ported memory controller that includes FPGA processing capabilities. Advantages include:

- Multiple I/O ports allow the controller to efficiently receive and transmit data on independent paths using the most practical protocol on each interlink.
- User Programmable Logic (UPL) within the controller can provide the processing power in an FPGA to offload the DSP processors.
- Large memory resources provide data buffering and temporary storage for intermediate results. Memory requirements for DSP processors can be reduced, and memory can be allocated to those processors in a manner that allows more efficient use of CPU cycles.
- Striding DMA that provides efficient corner turning to further offload processors.

FPGA-based System-on-Chip (SoC) products are starting to appear on the market and include features such as the multi-pointed controller shown in Figure 1.

User programmable logic and buffer memory

Many of the data transformations performed in DSP applications are highly repetitive and are efficiently implemented in FPGAs. Fast Fourier Transforms (FFTs), convolutions, FIR filters, and

IQ demodulation are among the processing often implemented in an FPGA. The parallel processing enabled by FPGA hardware implementations allows these operations to be performed much faster than using traditional processors, and the general-purpose or DSP processors can be offloaded to perform work that is not as well-suited to a hardware implementation – such as tasks that require decisions to be made based on the data. Fewer processors are thus needed in the system to complete all operations.

The large memory on a multi-ported controller can be used in many ways. Input or output data can be rate buffered. An input device may provide bursts of data, while real-time processing proceeds most efficiently at a steady pace. The buffered input data can be sent to processing units just in time for processing. Similarly, output buffering may be needed if the output device is not ready immediately as it comes from the real-time processing system.

The memory also may be needed between stages of processing. If the UPL is programmed to perform several independent kinds of processing such as FFTs or convolutions in several dimensions, large amounts of data may need to be stored between the stages. High bandwidth paths within the multi-ported controller allow the UPL to access the data several times without ever having to send the data over the external interlinks. The intermediate data can be stored in the buffer memory until enough accumulates to begin work on the next stage. Data sizes are application dependent and can vary significantly, from tens of megabytes to several gigabytes. The buffer memory also plays an important role in corner-turning operations.

Many FPGAs such as the Xilinx Virtex-II Pro and the Virtex-4 FX include embedded processors. A multi-ported controller based on these FPGAs may use these processors for additional data manipulation. The processors can also be used for control operations such as programming DMA and processing exceptions. This can offload the other processors in the system to help them concentrate on data manipulation and improve real-time performance.

Corner-turning DMA

A common problem in DSP applications is that data must be processed in multiple dimensions. In a 2-D image, for example, processing must first be performed on each row and then on each column of pixels. To accomplish this, the pixel matrix must be transposed. This is called *corner turning*.

On a multiprocessor system, distributed corner turning is needed to accomplish the transpose across all of the processors. Each row processor is assigned a subset of the rows. Each column processor is assigned a subset of the columns. When a row processor completes a row, it breaks up the row into pieces according to

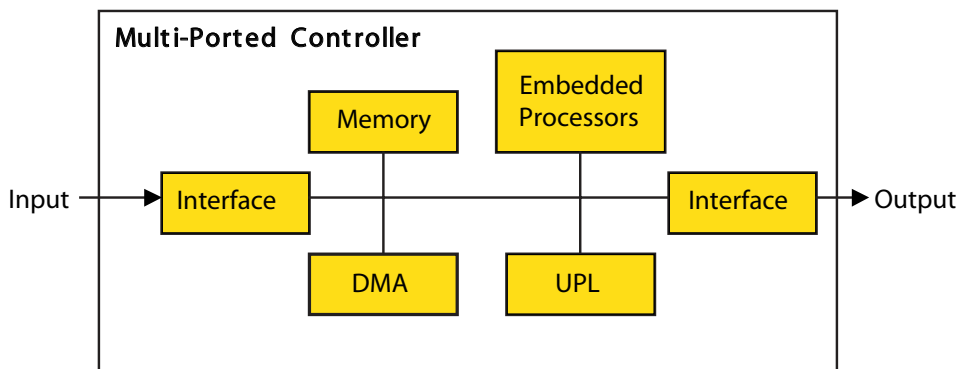


Figure 1

which column processor will work on each piece. It must then transmit each of these pieces to all of the column processors. The column processors receiving a piece of a row thus get just one element of each column. They must wait to accumulate enough rows before they can begin column processing.

This puts a large memory requirement on each column processor to buffer so many rows. Because each column processor must transpose the matrix it receives, many CPU cycles are also spent on corner turning. Figure 2 shows distributed corner turning, where data is distributed from the output of each row processor to every column processor by a chain of L DMA block transfers.

By adding a striding DMA that performs matrix transposition to the multi-ported controller, processors can be relieved of this duty.

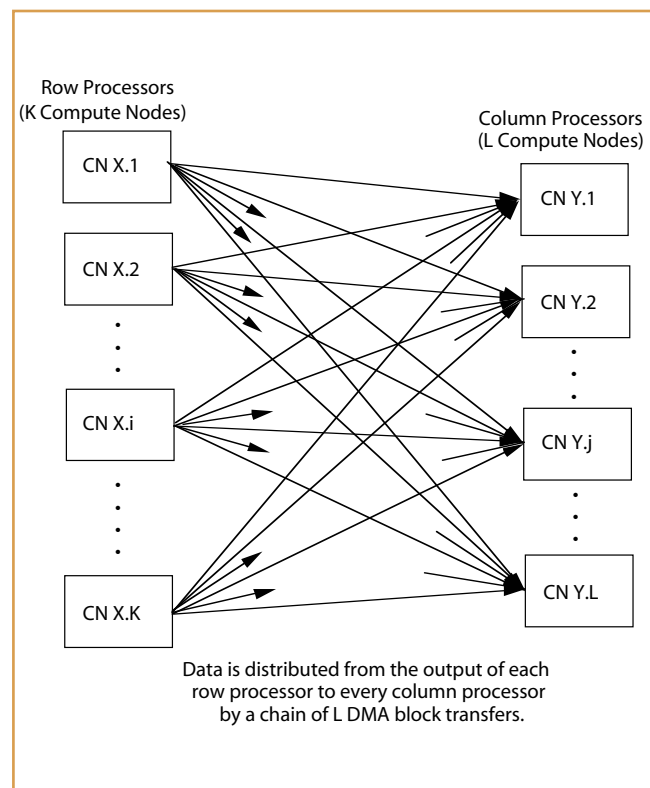


Figure 2

Since the processors can better concentrate on data transformations, fewer processors are needed to accomplish all the work. The savings on processors offset the cost of the multi-ported controller.

In such an application, the first stage of processing is performed by the FPGA, which works on rows of data received from the input device. The controller uses its large buffer memory to save the results of transformations on many rows as can be seen in the striding DMA in Figure 3. After completing the calculations on each single row of data, internal DMA built into the controller transfers the results of the calculation into buffer memory. These rows of results are collected to form a

MISSION CRITICAL

VME/cPCI
data storage modules



Extreme Comprehensiveness: We offer the most comprehensive VME/cPCI storage product line in the world, offering device alternatives for any standard or unique application.

- Solid State Disk • Removable Hard Disk
- Tape Drives • Optical Disk • PCMCIA Adapter

Extreme Performance: Our VME products feature extreme speed, capacity and ruggedly reliability with 320 MB/sec throughput enabled by LVD SCSI technology, storage capacity of more than 600 GBs per module and a 1,400,000 hour MTBF.

Extreme Quality: Phoenix International is the only manufacturer of VME data storage products that is ISO 9001:2000 Certified.



PHOENIX
INTERNATIONAL

Phoenix International Systems, Inc. An ISO 9001:2000 Certified SDVOSB

714-283-4800 • 800-203-4800 • www.phenxint.com




Figure 3

matrix in buffer memory. When enough rows have accumulated to allow efficient column processing, each column is transmitted to one of many DSP processors. The column consists of just one element, or a small number of elements, in each row.

The elements of the rows are stored consecutively in buffer memory. The consecutive elements of columns are separated in memory by the corresponding elements of all of the other columns, so they are not stored consecutively in memory. The DMA to the DSP processors must “stride” past all of these other columns. The striding DMA effectively performs a scatter-gather operation, collecting the nonconsecutive elements from the multi-ported controller’s buffer memory and transferring the elements into consecutive memory locations on the DSP processor.

Any interlink can support this striding DMA. Since the data is sent to the DSP processors in order of consecutive memory locations, it does not matter to the DSP processor or the interlink that the DMA engine in the multi-ported controller fetched the data from nonconsecutive locations in its own memory. This process implies, however, that the striding DMA must reside on the multi-ported controller.

All in stride

Figure 4 further illustrates this striding operation. The strides are local to the DMA master, so this is called *local striding*. The transferred segment size is the same on both the local system, the multi-ported controller, and the remote system. The multi-ported controller skips over the local stride during its access to its local memory. The multi-ported controller may transmit data to the DSP processors or receive data from them.

Figure 5 illustrates remote striding. The multi-ported controller is now shown on the right, but it is still acting as the DMA master. It must use separate transactions to access each segment on the remote system. Once again, it can transmit data or receive data. While not as efficient as local striding since separate transactions must be used for each segment, this can be useful if the remote system is not capable of mastering its own striding DMA.

Because the multi-ported controller buffers the data until a sufficient number of rows has been accumulated, the memory requirement of the column processors is greatly reduced. In some cases without a multi-ported controller, the column processors are memory-limited. Compromises have to be made to reduce the size of the rows or columns in order to keep the matrix size within the limits imposed by the system.

By using a multi-ported controller, larger rows or columns become practical because the column processors never have to deal with more than a few columns at a time. Since some operations, such as FFTs, require fewer instructions per pixel when performed

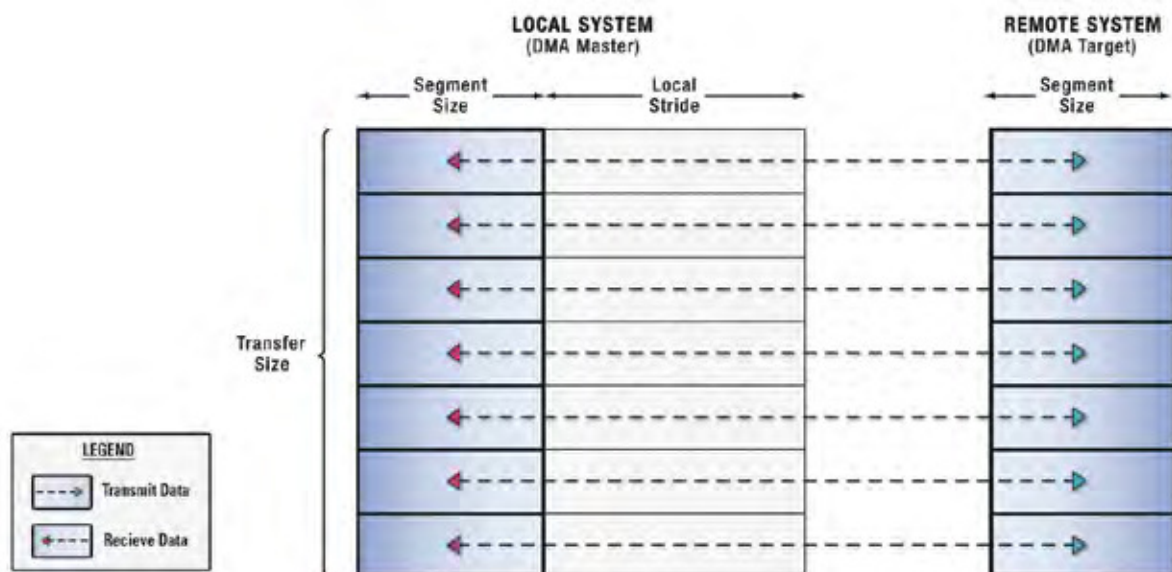


Figure 4

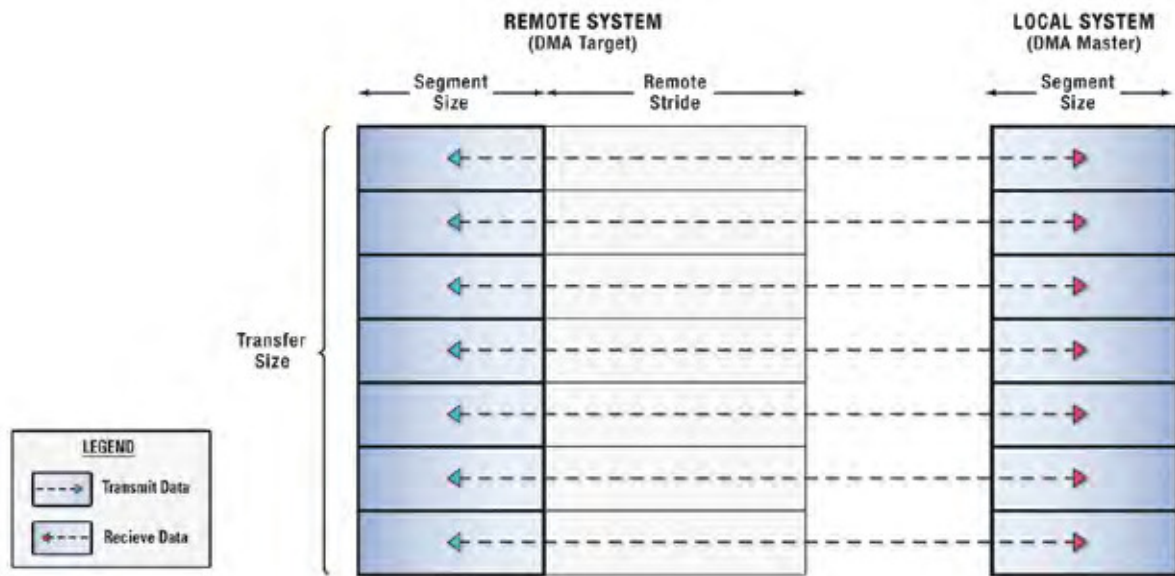


Figure 5

on larger columns, the CPU time needed to process the data can be reduced. The result is that fewer processors, each with less memory, will be able to complete the job.

Real-world controller

While FPGA processing can reduce the need for traditional DSP processors, a multi-ported controller tightly integrated with a UPL

can reduce that requirement even further. Having a multi-ported controller reduces the memory requirements for other processors, and the remaining memory can be used for longer columns. This allows for more efficient processing. Throughput is improved by using the most efficient protocol on each side of the multi-ported controller, without sacrificing efficiency to protocol translation, which eats CPU cycles.




Figure 6

Based on the Xilinx V-4 series of FPGAs, the CoSine SoC from Micro Memory combines UPL processing capabilities, a large multi-ported memory controller, and a striding DMA in a fully preconfigured solution (Figure 6). With this integrated functionality, CoSine provides a proven design and powerful platform for building DSP systems. +

Richard Mathews has been at Micro Memory for the past three years, where he currently leads Software Engineering. During that time, he has played key roles in both hardware and software system architecture. Prior to joining the company, Richard spent nine years at Sun Microsystems where he was a leading member in the Solaris port to x86, developed a course for writing Solaris device drivers, and worked with Tier I OEMs that utilized the ChorusOS. After studying physics at Caltech, Richard joined Locus Computing Corp. where he specialized in clustering mechanisms, memory management, load balancing, and remote file replication.



For more information, contact:

Micro Memory, LLC

9540 Vassar Ave.

Chatsworth, CA 91311

Tel: 818-998-0070

Fax: 818-998-4459

Email: sales@micromemory.com

Website: www.micromemory.com



Tapping into solid-state storage benefits for low-power and small form factor applications

By Gary Drossel

There are several advantages in using solid state technology in low-power and battery-powered applications intended for use in the military field. Enhanced security and power management features that are especially useful in mission-critical applications will be discussed, along with recent enhancements to solid state technology that have made these features possible.

Due to the harsh environments in which they are utilized and the sensitive data they contain, it has long been understood that storage solutions used in mission-critical military applications require ruggedness and advanced security beyond what is available in the average 2.5-inch rotating drive. For specialized field applications including wearable computers and mission data recorders, solid state technology is currently the most effective storage choice, increasing the reliability of storage exponentially in terms of both environmental viability and data security.

An evolution in solid state technology

Traditionally, solid-state storage solutions designed to satisfy these high-capacity and advanced data security requirements have been mechanically confined to 2.5-inch or 3.5-inch hard drive form factors. This is not only because of the number of solid state memory components required to achieve the desired capacity but also because of the physical size of the microprocessor and associated logic used to provide the host system interface and the solid-state memory management algorithms. These circuits have neither been able to scale to smaller form factors nor have they been able to achieve power consumption less than the typical 2.5 watts of rotating hard drives. This makes these technologies impractical for battery-operated or wearable computers.

New technology has been developed that has broken this paradigm by using ultra-low

power consuming proprietary Reduced Instruction Set Computing (RISC) technology to perform the same functions of host interface and solid-state storage management. The difference, however, is that this technology consumes less than one-tenth the power of the traditional solutions (0.2 watts typical) while achieving a mechanical footprint that allows traditional solid-state and rotating hard drives to fit in several industry standard form factors such as Type I CompactFlash, Type II PC Cards, and plug-in modules. Table 1 illustrates the mechanical advantages of this scalability:

Form Factor	Maximum Capacity	Dimensions	Weight L x W x H
2.5" Drive	32 GB	101 x 69.9 x 9.5	> 112 g
Plug-in Module	4 GB	50.8 x 5 x 22.9	9 g
Type II PC Card	32 GB	85.6 x 54 x 5	50 g
Type I CF	8 GB	36.4 x 42.8 x 3.3	15.2 g
Assumes no PCB or component stacking, which can dramatically reduce reliability in high-shock and vibration applications.			

Table 1

Enhanced data security

There are several levels of enhanced data security required by different types of military applications. These levels can range from *data sweep* – also known in the industry as *fast erase* – to die-level damage to the nonvolatile storage components. When initiated by the host, usually with a vendor-specific command, the data sweep function erases all the user data in seconds. The solid-state storage solution can then be reformatted and reused. This is defined as a *recoverable operation*.

The engineering trade-off of erase time versus power consumption must be considered in a data sweep application as shown in Figure 1. The more power that is available from the host, the more operations the controller can do in parallel, resulting in a faster sweep time. Contrast

the data in Figure 1 to a Flash card that can erase data at 10 MBps using standard ATA commands. The enhanced sweep technology erases the 2 GB solid-state drive in around 5 seconds, where it would take the Flash card approximately 3 minutes and 20 seconds to accomplish the same task. The advantages of this technology in combat situations are readily apparent.

Data purge technology takes the sweep technology one step further by also erasing all data control blocks and, if applicable, the firmware that comes with the controller for a particular drive.

This operation is defined as being *non-recoverable*; that is, the solid state drive cannot be reinitialized or reformatted. The purge function only adds a few milliseconds to the data erase time and virtually no current, so the choice between implementing a sweep function or a purge really becomes a matter of determining whether or not the drive is to be used again. Once the drive has been purged, no data remains and the drive is deemed secure.

Power management

Traditional solid-state drive architectures can in essence be thought of as a "PC in a drive form factor." These drives often employ a relatively high-end microprocessor, an FPGA for logic, solid-state memory components, and some type of RAM used as a buffer – hence the

VECTOR


ELECTRONICS & TECHNOLOGY, INC.
A FINE TECHNOLOGY GROUP



SERIES 790 RUGGEDIZED CHASSIS

BUILT FOR EXTREME AND
TOUGH CONDITIONS.

- Ⓢ WELL SUITED FOR MILITARY / AERONAUTICS
- Ⓢ VME64x, VME64 OR CPCI® COMPATIBLE
- Ⓢ MIL 461-D, E FOR SHOCK/VIBRATION, EMI/RFI
- Ⓢ CE AND FCC CLASS B, PART 15
- Ⓢ MANY POWER OPTIONS



790 CHASSIS



792 CHASSIS



ALSO FEATURING OUR
SERIES 800 TOWER ENCLOSURE
WITH REMOVABLE INTERNAL SUBRACK

VECTOR SPECIALIZES IN:

UNPARALLELED FACTORY SUPPORT
CUSTOM CONFIGURATIONS
SHORTER LEAD TIMES

(800)423-5659

WWW.VECTORELECT.COM



Made in U.S.A

RSC# 47 @www.mil-embedded.com/rsc




Figure 1

2.5 watt power requirement and the need for a larger mechanical footprint and (relative) lack of scalability.

The use of RAM, either DRAM or SRAM, as a buffer introduces the new design challenge of what happens to the integrity of the drive in the event of a power disturbance. RAM is a volatile memory architecture and if power is lost during a write operation, not only is the data lost, there is a good possibility that error correction and checksum of a particular sector won't match the data of that

buffer to be properly written to the non-volatile storage components and perform an orderly shutdown. It is important to note that this does not guarantee that the file being written won't be truncated. It may not be a useful file, but at least the drive is not corrupted.

An alternative solution is to design the solid state drive to utilize the smallest buffer possible. Only the highest end of data recording applications truly requires data rates in excess of about 15 MBps, so other more transactional types of

sector. The next time the controller goes to access that location, it sees the discrepancy, considers the sector bad, and may use one of a finite number of spares. If this happens too often, the drive can cease to operate because there are no spares remaining.

Solid-state drive designers attempt to overcome this scenario by integrating some type of optional power holdup circuit. This usually involves either a battery backup or large capacitor circuit, which can significantly increase the mechanical dimensions of the drive. The circuit is designed to allow the contents of the RAM


applications may be willing to sacrifice some speed for a high level of drive integrity. Additional reliability can be achieved by also integrating some voltage detection circuitry to allow the storage solution to make some very rapid decisions and perform an orderly shutdown.

The optimal solution is to integrate as much of this technology into the controller architecture itself so that mechanical footprints can be minimized. Figure 2 diagrams how a controller with voltage detection, regulation, and logic circuitry can be combined with a solid-state memory array to collectively maximize drive integrity.

More solid state drives for military

The application-enhancing ruggedness, scalability, advanced data security, and power management capabilities realized by incorporating solid state technology into military applications are still being discovered as designers replace existing mechanical and magnetic drives and add features to new designs. These new features cannot, however, come at the expense of increased size and power consumption – especially in wearable or battery-operated field computers. Therefore, storage technologies tailored to the military market must continue to evolve. ✚

Gary Drossel is responsible for managing marketing and business development for SiliconSystems' complete product line.



A 15-year industry veteran, he has played a leading role in developing the company's marketing strategy, including product roll-out and customer introduction. Drossel received a BS in electrical and computer engineering from the University of Wisconsin.

To learn more, contact Gary at:
SiliconSystems, Inc.

26940 Aliso Viejo Parkway
Aliso Viejo, CA 92656

Tel: 949-900-9400

Fax: 949-900-9500

E-mail: gdrossel@siliconsystems.com

Website: www.siliconsystems.com

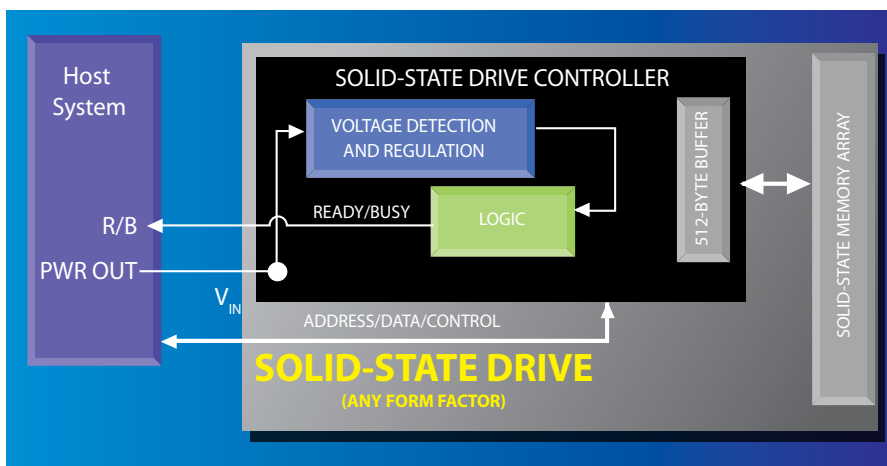


Figure 2

The Choice is Clear for Sonar, Seismic, and Recording Studio Applications

The Latest Data Acquisition Boards are more precise, use less power, less noise, software selectable I/O options, and are lower in price than the competitors. NOW your choice is easy. See the comparison table below and visit our website for full technical specs:

www.generalstandards.com/delta-sigma.php

Get our free report: **Top Ten Things to Remember Before Specifying a Data Acquisition Board.** Even if you don't buy from us, don't design in anyone's board until you have the answers indicated in this report.

The following is a comparison of Delta-Sigma Data Acquisition Board Characteristics of General Standards' offering and the leading competitor:

Key Features	General Standards	Leading Competitor
Power Dissipation/ 32-Channel	12 Watts Typical**	35 Watts
Phase Skew	55ns (0.1 deg at 5KHz)	Not specified
Crosstalk	96dB	90dB
SINAD	93dB (PCI, cPCI, PMC, & PC104+)	86dB (cPCI), 90dB (PCI)
Gain Accuracy	+/- 0.1mV, +/- 0.1 percent	Not specified
Sample Rate	200K per Chan (PCI, PMC, PC104+ & cPCI)	108K per Chan (PCI) 216K per Chan (cPCI)
Industrial Temp Range	-40° to 85°C	-40° to 85°C
Commercial Temp Range	0° to 65°C	0° to 50°C
Cost for 32-Channel	\$5,995 (Qty 1)*	\$9,000

* Cost is for 32-Channel Commercial Temp Version. Quantity discounts available. Conformal coating available.

**Power is for the 32-channel version of the board. Power dissipation for other versions is 6 Watts typical.



We also offer drivers, including
VxWorks, Linux, LabView, and
Windows

**Free Loaner
Board Offer!**

**Call for Details
(800) 653-9970**

General Standards

We ARE the Right Choice for Cutting Edge Data I/O







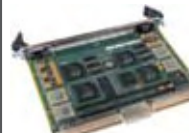





And You Can Call Us On It!

www.generalstandards.com

RSC# 49 @www.mil-embedded.com/rsc

(800) 653-9970













SBCs for DSP and reconfigurable computing

Company Name/Model Number		Form factor	Processor speed (Fastest available)	Processor manufacturer & model	
Blue Chip Technology www.bluechiptechnology.co.uk					
		ETX C3/C7	ETX	2 GHz	Via C3 C7
		MagnumX		2 GHz	Via C3 C7
Curtiss-Wright Controls Embedded Computing www.cwcembedded.com					
		Champ-AV IV	VME	6 GHz	Motorola PowerPC 7447/7448
		Manta QX3	VME	1 GHz	Motorola PowerPC 7457
		Champ-FX	VME	528 MHz	IBM 405 PowerPC

Product Selection Guide

Max volatile memory	Max nonvolatile memory	FPGA resources	Ports/slots					Conduction cooling available	
			Ethernet	Mezzanine	SCSI	Serial	USB		Other I/O
512 MB	1 GB	N/A	1	0	0	2	4	IR port, audio, parallel, LCD, CRT	Yes
512 MB	1 GB	N/A	1	0	0	2	4	Mini PCI, PCI, 16-channel GPIO, CRT, LCD, and LVDS	Yes
2 GB	256 MB	Core functionality through onboard timers	2	2	1	4	1		Yes
2 GB	1 GB	N/A	2	1	1	6	1		Yes
64 MB	N/A	Virtex-II Pro	1	2	1	8	1		Yes











Product Selection Guide

Company Name/Model Number			Form factor	Processor speed (Fastest available)	Processor manufacturer & model
Danville Signal Processing, Inc. www.danvillesignal.com					
		dspstak DSP boards		400 MHz	Analog Devices SHARC
Hunt Engineering www.hunt-dsp.com					
		HERON-USB-18		400 MHz	Xilinx Power PC and TI C6000
Intel Corporation www.intel.com					
		MPCBL0001	PMC	2 GHz	Intel MPCBL0001
		MPCBL0010	ATCA	2.8 GHz	Intel MPCBL0010
Tews Technologies www.tews.com					
		TVME8240	VME	250 MHz	Motorola MPC8240
		TVME8300	VME	300 MHz	Motorola MPC8245

Product Selection Guide

Max volatile memory	Max nonvolatile memory	FPGA resources	Ports/slots					Conduction cooling available	
			Ethernet	Mezzanine	SCSI	Serial	USB		Other I/O
8 MB	512 KB	Altera Cyclone FPGA available on some boards	0	4	0	1	2	SPI, I/O expansion port, JTAG, built-in ICE (USB)	No
288 MB	16 MB	Virtex-4 FX12, LX60/SX35 Virtex-II and Virtex-II Pro	0	4	0	3	1	210 Mhz ADC, 160 Mhz DAC, LVDS/LVTTL, FPDP, camera link	No
4 GB	266 MB	N/A	2	1	1	1	1	Dual Gigabit Ethernet, optional dual Fibre Channel	Yes
4 GB	1 MB	N/A	2	1	1	0	1	Dual GB Ethernet base and fabric interface	Yes
64 MB	8 MB	Optional IP FPGA mezzanine	1	4	1	2	0	Four industry pack sites	No
64 MB	8 MB	Optional PMC or IP FPGA mezzanine	1	4	1	2	0	Optional IP or PMC expansion	No

Product Selection Guide

Company Name/Model Number			Form factor	Processor speed (Fastest available)	Processor manufacturer & model
		TVME8400	VME	300 MHz	Motorola MPC8245
Thales Computers www.thalescomputers.com					
		PowerEngine7	VME	1 GHz	IBM PPC 750GX
VMETRO Inc. www.vmetro.com					
		PowerMIDAS C5000	CompactPCI	800 MHz	AMCC PPC440GX
		PowerMIDAS M5000	VME	667 MHz	AMCC PPC440GX
		VPF1	VME	1.25 GHz	Freescale 7447A
<p>1. Data is from the OpenSystems Publishing online database as of presstime. Vendors are encouraged to update and add listings to our online database. To update, see www.opensystems-publishing.com/vendors/submissions/np/.</p> <p>2. Entries have been edited per OpenSystems Publishing standards and are listed at the sole discretion of the editor. OSP offers no warranties, accepts no liability, nor offers indemnification from omissions or errors.</p>					

Product Selection Guide

Max volatile memory	Max nonvolatile memory	FPGA resources	Ports/slots					Conduction cooling available	
			Ethernet	Mezzanine	SCSI	Serial	USB		Other I/O
64 MB	8 MB	Optional PMC FPGA mezzanine	1	2	1	2	0	Dual PMC sites	No
512 MB	128 KB	N/A	2	2	1	6	1		Yes
1 GB	32 MB	N/A	2	2	1	2	1	Up to 2 x 2 Gbps optical fiber and PIM	No
256 MB	32 MB	N/A	3	2	1	0	1	2 x 2 Gbps optical fiber, RS-232/422, RACE++	No
512 MB	128 MB	2 V2Pro P70 FPGAs with dedicated QDR SRAM and SDRAM	2	1	1	4	1	Rocket I/O, direct FPGA	Yes

To view the full product selection guide, go to
www.mil-embedded.com/products/guides/reconfigurableSBCs

By Sharon Schnakenburg

Carrier board: PMC	56
Chips & Cores: FPGA	56
Connector: Other	56
Data acquisition	56
DSP resource boards: PCI-X	56
DSP resource boards: PCI/ISA	58
DSP resource boards: PMC	58
Fabrics: Fibre Channel	58
Fabrics: PICMG 2.16 Ethernet	58
Generators: Function/waveform	58
MIL-STD-1553	58
PCMCIA/PC Card	59
Processor: PowerPC	59
Routers/Switches	60
Server	60
Software-defined radio	60
Software: Modeling tool	60
Turnkey system	61
Video: Frame grabber	61

CARRIER BOARD: PMC

N.A.T. GmbH

Website: www.nateurope.com

Model: NcPCI-PMC RSC No: 21630



PCI interface and compliance: Intel i21555 PCI-to-PCI bridge (66 MHz, 64-bit), PCI Rev. 2.2 • H.110 bus: Agere T8110 TSI, H.110 on CompactPCI J4 connector, SCbus on PMC P14 connector • PMC slots: two 64-bit/66 MHz PCI Rev.2.2 IEEE P1386.1/Draft 2.4a compliant PMC slots on the PCI internal bus, I/O P14 connector used as SCSA bus connected to the T8110 TSI H.110 device • Full hot swap capability • Power consumption: 3.3 V, 0.5 A (typical), 5 V, 0.1 A (typical) • 0 °C to +60 °C operating temperature range, with forced air cooling • Standard compliance: PCI Rev. 2.2, PICMG 2.1 R2.0, PICMG 2.5 R1.0, IEEE P1386.1/Draft 2.4a

CHIPS & CORES: FPGA

SBS Technologies, Inc.

Website: www.sbs.com

Model: TS-CPCI-8001 RSC No: 21058



High-performance, FPGA-based computing platform for demanding signal and image processing applications • Suitable for SDR, sonar, and other high-performance signal processing applications • Dual Stratix EP1S80 FPGAs onboard • Two PMC sites with 1 Gbps high-speed data paths to each FPGA • 8 Gbps connection between FPGAs for a common processor fabric • Windows XP/2000, Linux, and Integrity software support

CONNECTOR: OTHER

Molex, Inc.

Website: www.molex.com

Model: Optical Military Cables RSC No: 21189



Military-style connector solutions available with fiber optic MIL-T-19504 approved termini, coupled with COTS cabling and fiber optic connectors • Suitable for a variety of military, avionics, flight control, mobile tactical field command platforms, EMI-sensitive, and security applications • Utilizes ruggedized connectors designed for harsh environments • Various finish types offered with plug, jam-nut, and wall-mount receptacle options • Can be terminated to both single-fiber and ribbon cable, available in 2- to 29-pin configurations to meet MIL-STD-1560 requirements • Custom lengths for COTS cabling and fiber routing options

DATA ACQUISITION

Great River Technology, Inc.

Website: www.greatrivertech.com

Model: Gravity 64/32 HOTLink II Data RSC No: 21376



Flexible, off-the-shelf solutions for high-speed, point-to-point data transfer using HOTLink II Fibre Channel • Supports baud rates from 200 Mbps to 1.5 Gbps • Variety of uses includes engineering development and factory test of HOTLink based interfaces • Can be used with almost any customer-defined HOTLink protocol, with data rates of 160 Mbps to 1.5 Gbps • Flexible connector options • Fiber option single or dual channel, FCN single or dual channel, or SMA • Flexible protocols • Users can select one of three modes: 4 byte packets, variable length packets, or fixed length packets

DSP RESOURCE BOARDS: PCI-X

Signatec, Inc.

Website: www.signatec.com

Model: PMP1000 RSC No: 21253



Parallel DSP Board • 72 GIPS peak processing power • Continuous input data processing up to 1000 MBps (depending on application) • Up to nine TI C6414 DSPs • 64 MB memory in each of eight main processing DSPs, totaling more than 512 MB DSP memory • Unique program execution processor for dynamic thread allocation • Advanced parallel DSP OS simplifies user programming • Customizable external interface for accommodating any type of user data channels • 2.5 Gbps of total external I/O via three interfaces • 2600 MBps internal I/O via switching network • 64-bit PCI-X local bus full-length board • Designed to create true high-speed, real-time systems in otherwise non-real-time environments (such as PCs)

WE'VE GOT YOUR BACK.



Ampro ETX-800

Ground Force Systems Specialists and communications experts must think clearly under fire, and respond with lightning speed. The rigorous physical and mental challenges of combat demand field equipment built to ensure mission success because the data captured, analyzed and transmitted helps make split-second decisions that save lives.

It's no wonder so many manufacturers choose Ampro SBCs and modules for rugged embedded applications. Ampro has always been the leader in modular embedded technology. And with their recent deployment of the first rugged ETX module, Ampro continues to out-rank all others in the field.

AMPRO. BECAUSE MORE THAN YOUR REPUTATION IS AT STAKE.

Call today (408) 360-0200 or visit: www.ampro.com/RuggedETX



Ampro is a registered trademark of Ampro Computers, Inc. All other trademarks are the property of their respective owners. © 2005 Ampro Computers, Inc. All rights reserved.

RSC# 57 @www.mil-embedded.com/rsc

DSP RESOURCE BOARDS: PCI/ISA

Innovative Integration

Website: www.innovative-dsp.com

Model: M6713

RSC No: 21886



DSP + FPGA board with novel architecture for advanced data capture and real-time control in PCI systems • Designed around Texas Instruments' floating-point DSP for high-speed, high dynamic range signal processing, and Xilinx's latest FPGA for unlimited customization of I/O peripherals • PCI 64-bit/66 MHz • Two OMNIBUS I/O module sites • Reconfigurable FPGA option up to 1.5 million gates • Supports multiple card I/O synchronization • Extensive software support in source form • Custom logic development supported for FPGA

DSP RESOURCE BOARDS: PMC

VMETRO

Website: www.vmetro.com

Model: PMC-FPGA03/F

RSC No: 21357



XC2VP50 Xilinx Virtex-II Pro FPGA PMC supporting two or four fiber optic I/O channels, each of which connects to a RocketIO channel on the FPGA • 64-bit user programmable data port (PMC user I/O P14) • Two banks DDR SDRAM (64 MB per bank) • Three banks QDR-II SRAM (up to 8 M x 18-bit per bank) • 2 or 4x fiber-optic transceivers (front panel) • Bandwidths of up to 3.125 GBps per port • 64-bit user programmable data port (PMC user I/O P14) • 4 MB Flash memory

FABRICS: FIBRE CHANNEL

Data Device Corp. (DDC)

Website: www.ddc-web.com

Model: FibreAccess

RSC No: 21225



Based on DDC-developed Intellectual Property, enabling long life cycle support and customizability for military applications • Dual-channel operation • True conduction cooled, flyable design at temperatures -40 °C to +85 °C and beyond • 1 or 2 GBps operation • Class 2 and 3 service including broadcast and multicast • 320 MBps throughput with memory-to-memory latency under 20 µs • Proven interoperability with both avionics-only and commercial fibre channel cards and switches • Rugged low-profile transceiver option offers approximately 50% of the footprint of a standard small form factor optical transceiver, enabling military/aerospace customers to minimize overall size and weight and to increase fiber optic cable bend radius

FABRICS: PICMG 2.16 ETHERNET

DSS Networks

Website: www.dssnetworks.com

Model: Metro-Switch Model 8261 RSC No: 21049



12-port PIGMG 2.16/VITA 31.1 6U gigabit switch fabric card • Unmanaged, lightly managed, or fully managed modes • Rugged versions available with extended temperature and conformal coating features for harsh environments • 850 nm multimode, 1310 nm single mode, and WDM fiber connector options • Simultaneous, fully independent operation on all ports, 32 Mbps-24 GBps • Fourth generation BCM5690/5695 switch fabric and BCM5464SRKB quad port transceivers from Broadcom • High-performance wire speed on all ports, 24 GB aggregate total, 1 MB onboard memory for packet buffering • Fully compliant to IEEE 802.3 specifications including auto negotiation • Available with OEM developer kit • VxWorks and Linux driver support

GENERATORS: FUNCTION/WAVEFORM

ZTEC Instruments, Inc.

Website: www.ztec-inc.com

Model: ZT530PCI

RSC No: 20700



A three-in-one modular instrument that provides a high-performance arbitrary waveform generator, a function generator with built-in waveforms, and a high-speed digital pattern generator • Paired analog signals with minimal channel-to-channel skew provide high-quality differential or IQ signal generation • Digital TTL outputs provide time markers, sync, or pattern stimulus • Instrument control achieved through command interface or register level access • Two channels simultaneous, 16-bit resolution • 160 MSps data rate with 2x, 4x, and 8x interpolation filter for up to 400 MSps • Up to 2 MS record length per channel • 32 TTL pattern outputs • Differential waveform synthesis

MIL-STD-1553

Ballard Technology

Website: www.ballardtech.com

Model: OmniBus Family

RSC No: 20634



Mixed protocols • High channel counts • User processor (PowerPC) • IRIG time and synchronization • Features one to four DSPs for protocol processing • Supports MIL-STD-1553, ARINC 429, ARINC 708, ARINC 717, serial ports (RS-422/232/423), CSDB, discrete I/O, and more • Available in several platforms (PCI, cPCI, VME, and USB/Ethernet) • Commercial temperature standard, industrial available • Modular design supports mixed-protocol applications • Can be upgraded/reconfigured

BMC Communications

Website: www.bmccorp.com

Model: PCI-UADI RSC No: 21267



A PCI universal avionics digital interface • 16-bit RISC low-power microcontroller that allows firm-ware programming in the field through a computer • Supports many communication protocols: MIL-STD-1553 dual redundant interface, ARINC 708/453, ARINC two-wire/six-wire protocols such as 429/575/571/572/581/582/615, and RS-232 • Error injection/detection • 32-bit message time tagging • Extensive BC and RT link list structures • Automatic/manual RT status bit • Advanced BC functionality • Onboard built-in test bus and hard-wired RT access • Fully compatible with the temperature, humidity, and vibration standards of MIL-STD-810E and EMI MIL-STD-416 • C Libraries, DLLs, Windows, and Linux device drivers

PCMCIA/PC CARD

Andor Design Corp

Website: www.andor-design.com

Model: PC401 RSC No: 22183




PCMCIA Type 2 to MIL-STD-1553 • Simultaneous BC, multiple RTs monitor • 16/32 bit time tag • Auto incrementing 128 k X 16 dual port static RAM Interface • Full error detection and generation • Onboard major minor frame timing • Program selectable protocol, A, B, custom • Programmable message parameters • Programmable output amplitude • Full/selective bus monitor • Full bus simulation/support validation/production test plans • One application for ISA, PCI, PCMCIA • Support for multiple assorted cards on the same platform

PROCESSOR: POWERPC

Creative Electronic Systems

Website: www.ces.ch

Model: RIOC 4070 RSC No: 21125



Conduction-cooled PowerPC 750Gx at maximum frequency • 512 MB global memory SDRAM at 800 MBps peak • CES-enhanced PowerPC-to-CompactPCI bridge • 16 independent linked list DMA channel engine • Two onboard PMC slots • Power-on/power-off control logic per PMC slot • High throughput DMA engine • 32 MB NOR with compressor • 256 MB NAND with high-speed file system • Multiple thermal sensors • Transparent multiprocessor extension with up to six MFCC 8446 companion modules • Extended BSPs for VxWorks 6.x and Integrity 5.x

EMBEDDED PLANET™

design.
The next
generation of
RFID
infrastructure.

develop.
Your products
based on our
platform.

deploy.
Your solution
faster.



Looking for Solutions to the RFID Equation?

PowerPC Computing Engine + Linux + PCMCIA RFID Reader



Quality is designed into every Embedded Planet product. Each product is designed to take you from prototype through production quantity volumes.

Superior **Documentation** ensures that you are up and running shortly after you open the box. Visit our web site to download our documentation.

Our engineers can **Support** your project with both hardware and software expertise to help you meet your time and budget constraints.

We can customize any of our modules for your application. Visit our website or contact us today for your complete solution.

4760 Richmond Rd / Cleveland, OH 44128
Tel: 216.245.4180 / Fax: 216.292.0561
www.embeddedplanet.com



RSC# 59 @www.mil-embedded.com/rsc

MILITARY EMBEDDED SYSTEMS October 2005 / 59

TURNKEY SYSTEM

Lanner Electronics, Inc.

Website: www.lannerinc.com

Model: EM-S300A

RSC No: 19111



Intel ULV Celeron CPU-based, wall-mountable mini workstation • Rugged construction for harsh environments • 10/100Base-T and AC'97 audio • Two COM ports, one parallel port, four USB ports, and one IrDA • Reserved CompactFlash Type I/II socket • One 2.5" HDD bay • Intel ultra-low-voltage Celeron processor at 650 MHz or 400 MHz (fanless) onboard • VIA chipset integrated with AGP 4X graphics • Supports 36 bits TFT/DSTN/STN LCD and two channels (2x18 bits) LVDS • One ATA/100 IDE channel • One built-in PC/104 expansion slot

VIDEO: FRAME GRABBER

ARVOO Imaging Products BV

Website: www.arvoo.com

Model: leonardo PMC64-CL

RSC No: 20721



PMC mezzanine video processor for CameraLink Base (24-bit) digital video • Onboard FPGA (Xilinx Virtex E 100/300) for preprocessing of acquired video data offers contrast stretching in gray value domain, RGB mosaic color restoration, and random 2D convolution filters • Onboard memory of 128 MB • PMC bus interface 64-bit/66 MHz maximum, resulting in an extreme transfer rate of 528 MBps • Software support for Windows, Linux, Real-Time Linux, QNX, and Solaris • Digital I/O lines available

When failure just isn't an option



- ✓ Radar Control
- ✓ Satellite Tracking
- ✓ High-altitude UAV
- ✓ Ship Guidance
- ✓ Global Positioning
- ✓ Missile Launch
- ✓ Airborne Surveillance



Hybricon performs

Hybricon Corporation is a world-class leader in the design and manufacture of electronic packaging solutions for high performance military and ruggedized COTS applications. Hybricon has a successful track record of proven and deployed COTS solutions that have been designed to meet stringent military requirements.

Call us at 1-877-HYBRICON (492-7426) or visit www.hybricon.com



Hybricon Corporation 12 Willow Road
Ayer, MA 01432 ISO 9001 Certified

www.hybricon.com

RSC# 6101 @www.mil-embedded.com/rsc

For more information, enter the product's RSC number at
www.mil-embedded.com/products

Crosshairs Editorial

By Chris A. Ciufu

We need SDR now



Recent worldwide events have shown that we can no longer wait for Software Defined Radios (SDRs) in military and civilian applications. For once, the obstacle is neither perception nor technology – it's simply a matter of money. And not that much money, since COTS equipment already exists. World events like last year's Indian Ocean tsunami, hurricanes Katrina and Rita in the United States, and the recent earthquake in Afghanistan's Muzaffarabad region and neighboring Pakistan all point to the need for interoperable communications equipment and infrastructure.

In emergency and chaotic situations, using portable wireless radio equipment is by far the most efficient way to communicate when disasters or terrorism wipe out land-based systems. Though robust, battery-backed up Plain Old Telephone Service (POTS) and cellular phone basestations are in fixed and often vulnerable areas subject to floods, earthquakes, or terrorism.

On June 4, Operation Atlas was the first airborne counterterrorism "war game" drill in the United States since 9/11. Nearly four months later, hurricanes Katrina and Rita struck. Both the Atlas exercise and the hurricanes uncovered critical and deadly weaknesses in the communications abilities of first responder civilian and military personnel. In short: Inadequate infrastructure, planning, and radios that don't talk between organizations lead to simulated – as well as actual – life and death situations.

Articles published in this issue of *Military Embedded Systems* by industry experts such as Spectrum Signal Processing, General Dynamics C4 Systems, Composable Logic, and SCA Technica (page 28) all define SDR as digital radios that can create multiple waveforms to enable interoperability between legacy radios. SDR has grown into the basis for the DoD's massive Joint Tactical Radio System (JTRS) program that seeks to bridge communications between dozens of legacy military radios and ultimately replace them.

The technology exists today, and JTRS progress remains strong with tens of billions of dollars in congressional funding (\$12-15 billion, according to my estimates, Boeing, and Xilinx). But the commercial and civilian worlds, including police, fire, city, and state officials, FEMA, and the Red Cross, lag far behind with outdated and RF communications that generally can't talk to each other, much less to the military such as the National Guard troops who eventually descended upon New Orleans. The fundamental technologies needed for SDR include low-power processors such as those from ARM used in cell phones, DSPs from companies like TI and Analog Devices, and partially reconfigurable FPGAs from the likes of Altera, Xilinx, and others. The software "middleware" called the SCA that's necessary to make it work has been developed by the DoD and is available from various sources and administered by the Software Defined Radio Forum (SDRF).

The cost of *failing* to deploy SDRs is well-documented. The war game Operation Atlas showed that only United Airlines and the FBI could communicate with the "hijacked" airplane via the traditional Aircraft Communications Addressing and Reporting System (ACARS). But officials participating in the exercise from Massachusetts' Massport, state police, Federal Air Marshal Service, Homeland Security, Coast Guard, FAA, and others could only wait for information to trickle in. A key failure noted at the end of the exercise: Information sharing and direct communications between agencies is lacking.

So, too, were the lessons from Hurricane Katrina. Do a Google search on "Katrina interoperable communications" and 33,500 hits are returned. With everything under water, only walkie-talkies and a few police networks survived. In one instance, FedEx noticed that one of their own radio antennas atop a 54-story building had survived. FedEx technicians and the Army's 82nd Airborne rigged a generator to the transceiver and distributed 125 walkie-talkies to emergency responders. Cell phone companies deployed portable base stations to replace damaged towers. The myriad agencies involved in disaster relief not only had no communications infrastructure, they couldn't have talked to each other even if they did.

SDR pioneers such as Vanu already sell and have deployed COTS systems such as the Anywave Base Station that can realize multiple cell phone protocols so that one basestation can talk to all major carriers via GSM, CDMA, WCDMA, and so on. The product is being tested at the Idaho National Laboratory. Vanu has an SDR version that mounts in a first responder's vehicle that can be programmed to bridge police frequencies with fire and other officials, allowing *different* radios to interoperate. And the equipment is based on low-cost Pentium processors and FPGAs.

Clearly the issue isn't lack of demand or technology. Congress recognizes that interoperability failures demand "immediate action." A report written by senators McCain (AZ) and Lieberman (CT) and representatives Harman (CA) and Weldon (PA) published in the Washington Post on 9/18 indicates we're no better off four years after 9/11, and that "Policymakers at all levels have failed...by not taking action on interoperability." The very next day, the Senate failed to agree on waiving budget rules that would provide funding for interoperable communications equipment. The bill failed 40-58.

As I write this, Hurricane Wilma is at Category 5 and is bearing down on Florida. Can we afford to wait another four years for SDR's promise of interoperable communications?



Chris A. Ciufu
Group Editorial Director

Extreme Engineering Solutions

Website: www.xes-inc.com

Model: XCalibur1002 RSC No: 22443



A 6U CompactPCI ruggedized, extended temperature blade • IBM 750GX PowerPC processor • Ruggedized and operational from -40 °C to +85 °C • Dual PCI-X PrPMC slots • Up to 1 GB 266 MHz DDR SDRAM • 4-80 MB soldered Flash and CompactFlash socket • Dual PCI-X PrPMC slots • Front panel serial and 10/100/1000 Mbps Ethernet ports • Two 10/100/1000 Mbps PICMG 2.16 backplane Ethernet ports • Complies with PICMG 2.0, 2.1, 2.3, 2.9, 2.16 • VxWorks, Linux, and QNX Neutrino V6.3 support

ROUTERS/SWITCHES

ACT/Technico

Website: www.acttechnico.com

Model: 6U GigE Switch RSC No: 22362



24 10/100/1000Base-T Ethernet ports • Switching capacity of 24 M packets (16 ports) and 36 M packets (24 ports) • Full-line rate switching engine • 8,000 MAC addresses • Port and MAC access control • Flow control and back pressure • Automatic or controlled learning and aging routing table • Link aggregation/port trunking • Jumbo frame of up to 10 KB • 256 active VLANs • Enhanced port mirroring • STP/RSTP algorithm for more reliable network communications • Four priority queues per port • Easy software updating via FTP download to local in Flash memory • Comprehensive power on built-in test • Thermal monitoring • Extended temp and rugged versions available

SERVER

NextCom

Website: www.nextcomputing.com

Model: Flextrex RSC No: 22828



A briefcase-sized portable server and mobile workstation • Open standards architecture • Processor configurations: single, dual, or dual-core Opteron processors • Multiple high-performance memory configurations • Multiple Gigabit Ethernet ports • Multiple SATA and USB 2.0 ports • PCI Express X16 expansion slot • Optional Ultra160 SCSI or fibre channel • Dual head graphics support • Multiple removable or fixed HDD configurations • Supports several operating systems: Linux, Microsoft Windows 2003 Server, XP Pro, 2000 Pro, Solaris x86, Trusted Solaris x86, Solaris 10, Win-UX concurrent use Linux and Windows, and multiple dual boot environments

SOFTWARE-DEFINED RADIO

Interactive Circuits & Systems Ltd. (ICS)

Website: www.ics-ltd.com

Model: PMC571 RSC No: 21093



4 million gate FPGA, providing the platform for application-specific software development within a Xilinx Virtex II environment • Single-channel, software radio rugged mezzanine • Sample rates up to 80 MHz and processing bandwidth of 40 MHz • 14-bit data conversion on ADC and DAC functions • Full 64-bit, 66 MHz PCI interface and a 64-bit general-purpose I/O • PCI 2.2 compliant (64-bit, 66 MHz), Master/Target Burst Mode (DMA) capable, with DMA chaining (scatter-gather) • Keyed for universal signaling • Sampling/conversion occurs following rising or falling edge of external trigger (programmable) • Air- and conduction-cooled ruggedization levels and extensive application and technical support data available

Spectrum Signal Processing

Website: www.spectrumsignal.com

Model: SDR-3000 RSC No: 22214



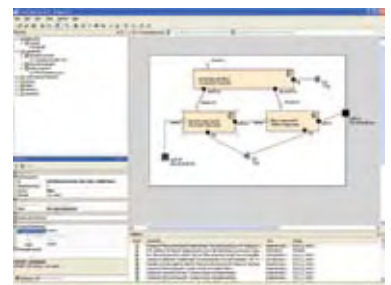
Software-Defined Radio platform • CompactPCI-based architecture • Xilinx Virtex-II FPGAs • PowerPC MPC7410 G4s • Optional TI TMS320C64X DSPs • Supports IF sampling rates up to 212 MHz • Supports 2-4 ADC channels and 2-4 DAC channels per slot • FlexFabric - serial rapid IO-based switched fabric connects all boards with deterministic, low-latency 320 MBps data paths • Supports up to 1,000 simultaneous transmit and receive communication channels per chassis • Includes optional SCA core framework • Supplied with full CORBA support (TAO) • Supports the VxWorks RTOS on every PowerPC node

SOFTWARE: MODELING TOOL

Zeligsoft

Website: www.zeligsoft.com

Model: Component Enabler RSC No: 20776



Creates visual models of SDR systems that adhere to the SCA standard as defined by the JTRS • Validates the models for compliance to the SCA standard • Generates correct-by-construction XML descriptor files and user-customizable documentation • Imports existing or third-party descriptor files, validates for standards compliance, visualizes, and integrates into a full system • Provides development team with a UML 2.0-based visual representation of system architecture • Provides early validation and generation of SCA 2.2, 2.2.1, and 3.0 compliant descriptor files • Improves overall quality and readability of waveform applications • Fits within current development processes




Battle Grade Ethernet™

SBS knows rugged Ethernet. With failover and BIT, our switches are battle ready.

OUR GIGABIT ETHERNET SWITCHES have been through boot camp—they've been tested, qualified, toughened up and generally put through the wringer. What emerges from this brutal test of character is Battle Grade Ethernet™.

These switches come with failover software and built in test (BIT). They are available in air or conduction cooled versions, with up to 24 ports for 6U or 3U CompactPCI® systems. Our most rugged switches can operate at temperatures



CP3-GEW12M3
3U CPCI 10-Port Layer-2 &
Layer-3 Managed Switch

from -40° to +85° C and have been designed to meet VITA 47 classes ECC1, ECC2, ECC3 and ECC4.

Ethernet is a proven technology in many military programs, and it is an important element of most networking applications. Which means Ethernet will have the kind of long term support that is critical in COTS programs. So if you're looking for Ethernet that has earned its stripes, give us a call. Sir.



SBS knows.

Find the Ethernet you're looking for at www.defense.sbs.com or call **800.SBS.EMDEDED**

RSC# 63 @www.mil-embedded.com/rsc



POWER OF TEAMWORK

At Curtiss-Wright Controls Embedded Computing (CWCEC), we understand that assembling the most advanced components and integrated systems is only part of the solution. That's why we design and build products that not only work well, but work well together. From host CPU board to mezzanine module to ATR chassis, we've made interoperability of critical system components our top priority.

That's how CWCEC delivers the power of teamwork. That's what makes us your best partner for your embedded solutions needs.

And that's Innovation in Motion.



Dy 4 Systems | Peritek | Primographics | Synergy | Systron | VISTA Controls

**CURTISS
WRIGHT** Controls
Embedded Computing